




RS/6000 Systems Handbook

2000 Edition

The ideal deskside reference for the latest RS/6000 models and features

Hundreds of tables and figures to accelerate your research

A required reading for all RS/6000 and AIX professionals



Jana Babnik-Gomiscek
Volker Haug
Jeanine Indest
Stephen Lutz
Shyam Manohar
Irene D. Sideris
Scott Vetter

ibm.com/redbooks

Redbooks



International Technical Support Organization

**RS/6000 Systems Handbook
2000 Edition**

August 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix K, "Special Notices" on page 773.

Second Edition (August 2000)

This edition applies to IBM RS/6000 Models 140, 150, 170, B50, 260, 270, F50, F80, H50, H70, H80, M80, S7A, S80, SP, and NUMA-Q. Related software offerings include AIX Version 4.3, program number 5754-C34, and subsequent releases.

This document created or updated on August 9, 2000.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JN9B Building 003 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 1999, 2000. 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	xvii
Tables	xxi
Preface	xxvii
The Team That Wrote This Redbook	xxvii
Comments Welcome	xxxi
Chapter 1. RS/6000 Introduction	1
1.1 RS/6000 History	1
1.2 RS/6000 Design Overview	2
1.2.1 Computer System Block Diagram Explanation	2
1.2.2 RS/6000 Microprocessor Architectures	7
1.2.3 RS64 and RS64 II Microprocessors	9
1.2.4 RS64 III Processor	10
1.3 RS/6000 New Technologies and Directions	13
1.3.1 64-Bit Technology	13
1.3.2 Processor Technologies	14
1.3.3 IBM POWER4 Microprocessor	14
1.3.4 Project Monterey	15
Chapter 2. RS/6000 Facts and Features Summary	17
Chapter 3. Workstations and Workgroup Servers	25
3.1 Model History and Uses	25
3.1.1 43P Model 140	26
3.1.2 43P Model 150	26
3.1.3 44P Model 170	26
3.1.4 43P Model 260	26
3.1.5 44P Model 270	27
3.1.6 Model B50	27
3.1.7 Upgrade Paths	27
3.2 RS/6000 43P 7043 Model 140 Overview	28
3.2.1 43P Model 140 Standard Configuration	28
3.2.2 43P Model 140 System Expansion	29
3.2.3 43P Model 140 Optional Features	30
3.2.4 43P Model 140 Configuration Notes	34
3.2.5 43P Model 140 Publications	34
3.3 RS/6000 43P 7043 Model 150 Overview	35
3.3.1 43P Model 150 Minimum Configuration	35
3.3.2 43P Model 150 System Expansion	36

3.3.3	43P Model 150 Optional Features	36
3.3.4	43P Model 150 Configuration Notes	39
3.3.5	43P Model 150 Publications	40
3.4	RS/6000 44P 7044 Model 170 Overview	40
3.4.1	44P Model 170 Minimum Configuration	41
3.4.2	44P Model 170 System Expansion	41
3.4.3	44P Model 170 Optional Features	42
3.4.4	44P Model 170 Configuration Notes	45
3.4.5	44P Model 170 Publications	46
3.5	RS/6000 43P 7043 Model 260 Overview	47
3.5.1	43P Model 260 Minimum Configuration	47
3.5.2	43P Model 260 System Expansion	48
3.5.3	43P Model 260 Supported Optional Features	48
3.5.4	43P Model 260 Configuration Notes	52
3.5.5	43P Model 260 Publications	52
3.6	RS/6000 44P 7044 Model 270 Overview	53
3.6.1	44P Model 270 Minimum Configuration	53
3.6.2	44P Model 270 System Expansion	54
3.6.3	44P Model 270 Supported Optional Features	55
3.6.4	44P Model 270 Dynamic Processor Deallocation	58
3.6.5	44P Model 270 Configuration Notes	59
3.6.6	44P Model 270 Publications	60
3.7	RS/6000 7046 Model B50 Overview	61
3.7.1	Model B50 Minimum Configuration	61
3.7.2	Model B50 System Expansion	62
3.7.3	Model B50 Optional Features	62
3.7.4	Model B50 Racks	65
3.7.5	Model B50 Configuration Notes	66
3.7.6	Model B50 Publications	66
	Chapter 4. Midrange Enterprise Servers	69
4.1	Model History and Uses	69
4.1.1	The Deskside Servers	69
4.1.2	The Rack-Mounted Servers	70
4.1.3	Upgrade Paths	73
4.2	RS/6000 7025 Model F50 Overview	74
4.2.1	Model F50 Standard Features	74
4.2.2	Model F50 System Expansion	77
4.2.3	Model F50 Optional Features	77
4.2.4	Model F50 Configuration Notes	84
4.2.5	Model F50 Publications	84
4.3	RS/6000 7025 Model F80 Overview	85
4.3.1	Model F80 Standard Features	85

4.3.2	Model F80 Optional Features	89
4.3.3	Model F80 Configuration Notes	98
4.3.4	Model F50 Upgrades	98
4.3.5	Model F80 Publications	99
4.4	High-Availability Cluster Server Solution HA-F80	99
4.5	RS/6000 7026 Model H50 Overview	101
4.5.1	Model H50 Standard Configuration	101
4.5.2	Model H50 System Expansion	103
4.5.3	Model H50 Optional Features	104
4.5.4	Model H50 Configuration Notes	109
4.5.5	Model H50 Publications	109
4.6	RS/6000 7026 Model H70 Overview	110
4.6.1	Model H70 Minimum Configuration	111
4.6.2	Model H70 System Expansion	112
4.6.3	Model H70 Optional Features	112
4.6.4	Model H70 RAS	119
4.7	High-Availability Cluster Server Solution HA-H70	120
4.7.1	Solution HA-H70 Overview	120
4.7.2	Solution HA-H70 Configuration Notes	121
4.8	RS/6000 7026 Model H80 Overview	122
4.8.1	Model H80 Minimum Configuration	123
4.8.2	Model H80 Optional Features	126
4.8.3	Model H80 Configuration Notes	133
4.8.4	Model H70 Upgrades	134
4.8.5	Model H80 Publications	134
4.9	High-Availability Cluster Server Solution HA-H80	135
4.9.1	Solution HA-H80 Introduction	135
4.10	RS/6000 7026 Model M80 Overview	136
4.10.1	Model M80 Product Positioning	137
4.10.2	Model M80 Description	137
4.10.3	7026 Model H80 To Model M80 Conversion	139
4.10.4	RS/6000 Model M80 Standard Configuration	141
4.10.5	Model M80 Optional Features	142
4.11	High-Availability Cluster Server Solution HA-M80	150
	Chapter 5. High-End Enterprise Servers	153
5.1	Introduction to the Servers	153
5.1.1	Model S70	153
5.1.2	Model S70 Advanced	154
5.1.3	Model S80	155
5.1.4	Model S80 Upgrade Paths	157
5.2	RS/6000 Model S70 Advanced Product Description	157
5.2.1	Model S70 Advanced Overview	157

5.2.2	Model S70 Advanced Minimum Configuration	166
5.2.3	Model S70 Advanced System Expansion	167
5.2.4	Model S70 Advanced Optional Features	167
5.2.5	Model S70 Advanced Configuration Notes	171
5.2.6	Model S70 Advanced Publications	172
5.3	RS/6000 Model S80 Product Description	174
5.3.1	Model S80 Processors	174
5.3.2	Model S80 Power Supply	177
5.3.3	Model S80 Memory	177
5.3.4	Model S80 Cabling	182
5.3.5	Model S80 I/O Bridge Bus	189
5.3.6	Model S80 PCI Buses	189
5.3.7	Model S80 Firsts	189
5.3.8	Model S80 Configurations	190
5.3.9	Model S80 Additional Features	191
5.3.10	Model S80 Remote I/O Subsystem Description	194
5.3.11	Model S80 Service Tools	198
5.3.12	Upgrade Paths to Model S80	199
5.3.13	Model S80 Configuration Notes	203
5.3.14	Model S80 Publications	205
5.4	SP External Node Attach and the S-Series	205
5.5	High-Availability Cluster Server Solution HA-S70 Advanced	208
5.5.1	Solution HA-S70 Advanced Value Proposition	208
5.5.2	Solution HA-S70 Advanced Configurations	209
5.6	High-Availability Cluster Server Solution HA-S80	211
5.6.1	Solution HA-S80 Value Proposition	211
5.6.2	Solution HA-S80 Minimum Configuration	212
5.6.3	Solution HA-S80 Highlights	213
5.6.4	Solution HA-S80 System Expansion	216
5.7	System Combinations in the S80 I/O Rack	217
Chapter 6. The RS/6000 Server and Solution Racks		223
6.1	Model 7015-R00 Overview	223
6.1.1	Model 7015-R00 Physical Specifications	224
6.1.2	Model 7015-R00 Configuration Suggestions	225
6.1.3	Model 7015-R00 Publications	226
6.2	Model 7014-S00 Overview	226
6.2.1	Model 7014-S00 Physical Specifications	227
6.2.2	Model 7014-S00 Configuration Suggestions	228
6.2.3	Model 7014-S00 Publications	229
6.3	Model 7014-T00 Overview	230
6.3.1	Model 7014-T00 Highlights	231
6.3.2	Model 7014-T00 Physical Specifications	233

6.3.3	Model 7014-T00 Configuration Suggestions	234
6.3.4	Model 7014-T00 Publications	235
6.4	Model 7014-T42 Overview	235
6.4.1	Model 7014-T42 Highlights	236
6.4.2	Model 7014-T42 Physical Specifications	239
6.4.3	Model 7014-T42 Configuration Suggestions	240
6.4.4	Model 7014-T42 Publications	241
Chapter 7. Large Scale Servers - RS/6000 SP Systems		243
7.1	SP Origins	244
7.2	SP Uses and Introduction	245
7.3	Information Sources for the SP	246
7.4	SP Hardware Components	246
7.5	The Frame	247
7.5.1	SP Frame Description	247
7.5.2	Contents of the Frame	247
7.5.3	SP Legacy Models and Differences of the New Frames	248
7.5.4	SP Legacy Models by Frame and Switch Types	248
7.5.5	SP Hardware Control and Supervision	249
7.5.6	SP Frame Dimensions	250
7.5.7	Frame Power Supplies	250
7.5.8	SP Frame Electrical Power Requirements	251
7.6	Internal Processor Nodes	253
7.6.1	Node Selection	254
7.6.2	SP Nodes at a Glance	255
7.6.3	High Nodes	259
7.6.4	Wide Nodes	265
7.6.5	Thin Nodes	275
7.6.6	332 MHz SMP Thin Nodes	281
7.7	The SP Switch	286
7.7.1	SP Switch Components	286
7.7.2	SP Switch Types	292
7.8	Control Workstation	293
7.8.1	High Availability Control Workstation	294
7.8.2	Supported and Unsupported SP Control Workstations	294
7.8.3	CWS Disk and Memory Requirements	297
7.9	The SP Switch Routers	298
7.9.1	Gigabit Ethernet Media Cards for the SP Switch Router	299
7.9.2	The Advantages of a Switch Router	300
7.10	SP Attached Servers - The S-Series	301
7.10.1	Floor Placement Considerations	301
7.10.2	HACWS Considerations with SP-Attached Servers	302
7.10.3	Pre-Install Required Meeting - System Assurance Reviews	304

7.10.4	64-bit Application Tolerance	304
7.10.5	Service	304
7.10.6	Ordering Through Joint Configurator Sessions	305
7.11	I/O Adapters Supported on PCI SMP Nodes	307
7.12	Supported Levels of AIX and PSSP on SP Nodes	309
7.12.1	Withdrawn Nodes and Their RS/6000 Equivalents	310
7.13	High Availability	311
7.14	System Partitioning	313
7.15	Peripheral Devices	313
Chapter 8. IBM NUMA-Q		315
8.1	NUMA Introduction	315
8.1.1	IBM NUMA-Q Architecture	315
8.1.2	IQ-Link	316
8.1.3	NUMA-Q SAN	316
8.1.4	NUMACenter	317
8.1.5	A Defining Technology	317
8.2	NUMA-Q Products, Solutions, and Services	317
8.2.1	NUMA-Q 2000 Enterprise Server	317
8.2.2	DYNIX/ptx	317
8.2.3	Specifications of NUMA-Q 2000 Enterprise Server	318
Chapter 9. Communications and Storage I/O Adapters		323
9.1	ISDN Adapter	325
9.1.1	Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)	325
9.2	System Adapters	325
9.2.1	IBM Short-wave Serial HIPPI PCI Adapter (# 2732)	325
9.2.2	IBM Long-wave Serial HIPPI PCI Adapter (# 2733)	326
9.2.3	S/390 ESCON Channel Adapter (# 2751)	326
9.2.4	IBM RS/6000 SP System Attachment Adapter (# 8396)	328
9.3	FDDI Adapters	329
9.3.1	SysKonnnect SK-NET FDDI-LP SAS PCI Adapter (# 2741)	329
9.3.2	FDDI SK-NET LP DAS Adapter (# 2742)	330
9.3.3	SysKonnnect SK-NET FDDI-UP SAS PCI Adapter (# 2743)	331
9.4	Asynchronous Adapters	332
9.4.1	8-Port Asynchronous Adapter EIA 232/RS-422 (# 2943)	332
9.4.2	128-Port Asynchronous Adapter EIA-232 (# 2944)	333
9.5	ARTIC Adapters	335
9.5.1	ARTIC960Hx 4-Port Selectable Adapter (# 2947)	335
9.5.2	IBM ARTIC960Hx 4-Port T1/E1 Adapter (# 2948)	336
9.5.3	ARTIC960RxD Quad Digital Trunk Adapter (# 6310)	337
9.6	WAN Adapters	338
9.6.1	2-Port Multiprotocol X.25 Adapter (# 2962)	338

9.7	Asynchronous Transfer Mode Adapters	340
9.7.1	Turboways 155 PCI UTP ATM Adapter (# 2963)	340
9.7.2	Turboways 155 PCI MMF ATM Adapter (# 2988)	341
9.7.3	Turboways 25 ATM Adapter (# 2998)	342
9.8	Ethernet and Token-Ring Adapters	342
9.8.1	10/100 Ethernet 10BaseTx PCI Adapter (# 2968)	342
9.8.2	Gigabit Ethernet - SX Adapter (# 2969)	344
9.8.3	PCI Ethernet BNC/RJ-45 Adapter (# 2985)	345
9.8.4	PCI Ethernet AUI/RJ-45 Adapter (# 2987)	346
9.8.5	4-Port 10/100 Base-Tx Ethernet PCI Adapter (# 4951)	347
9.8.6	4/16 Mbps Token-Ring Adapter (# 4959)	348
9.9	PCI Storage Adapters	349
9.9.1	SCSI-2 Fast/Wide RAID Adapter (# 2493)	349
9.9.2	PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)	350
9.9.3	PCI Universal Differential Ultra SCSI Adapter (# 6204)	351
9.9.4	PCI Dual Channel Ultra2 SCSI Adapter (# 6205)	352
9.9.5	PCI Single-Ended Ultra SCSI Adapter (# 6206)	353
9.9.6	PCI Differential Ultra SCSI Adapter (# 6207)	354
9.9.7	SCSI-2 Single-ended Fast/Wide Adapter (# 6208)	355
9.9.8	SCSI-2 Differential Fast/Wide Adapter (# 6209)	356
9.9.9	SSA RAID 5 Adapter (# 6215), SSA Fast-Write (# 6222)	357
9.9.10	Advanced SerialRAID Plus Adapter (# 6230)	358
9.9.11	Gigabit Fibre Channel Adapter (# 6227)	363
9.10	Hot-Pluggable Adapters	364
Chapter 10. Graphics Accelerators		367
10.1	Available Graphics Accelerators	368
10.1.1	POWER GXT3000P Graphics Accelerator (# 2825)	369
10.1.2	POWER GXT2000P Graphics Accelerator (# 2823)	373
10.1.3	POWER GXT300P Graphics Accelerator (# 2841)	376
10.1.4	POWER GXT250P (# 2851) and GXT255P (# 2852)	379
10.1.5	POWER GXT130P Graphics Accelerator (# 2830)	381
10.1.6	POWER GXT120P Graphics Accelerator (# 2838)	384
10.1.7	PCI Ultimedia Video Capture Adapter/S (# 2639)	386
10.1.8	Percentage Performance Improvement.	386
10.2	Withdrawn Graphics Accelerators	387
10.3	Limitations	389
Chapter 11. SCSI, SSA, and Fibre Channel Overview		391
11.1	SCSI	391
11.1.1	SCSI-I	391
11.1.2	SCSI-II	392
11.1.3	SCSI-III	392

11.1.4	SCSI Repeaters	399
11.1.5	Summary of SCSI Specifications	400
11.1.6	General SCSI Cabling Considerations	400
11.2	Serial Storage Architecture Overview	403
11.3	Comparison between SCSI and SSA Architecture	404
11.4	Fiber Channel-Arbitrated Loop (FC-AL)	405
Chapter 12. Internal Storage Devices		407
12.1	Disk Drives	407
12.1.1	9.1 GB Ultra-SCSI 1" High Disk Drive (# 2502)	408
12.1.2	18.2 GB Ultra-SCSI 1" High Disk (# 2503)	408
12.1.3	9.1 GB 10K RPM Ultra-SCSI 16-Bit Enhanced Disk (# 2906)	409
12.1.4	9.1 GB Ultra-SCSI 16-Bit 1" High Disk (# 2908)	409
12.1.5	18.2 GB Ultra-SCSI 1" High Disk (# 2909)	410
12.1.6	9.1 GB Ultra-SCSI 16-Bit Hot Swap Disk Drive (# 2913)	410
12.1.7	9.1 GB 10K RPM Ultra SCSI 1" High Hot Swap Disk (# 3002)	411
12.1.8	9.1 GB Ultra2 SCSI 16-Bit 1" High Disk (# 3021)	411
12.1.9	18.2 GB Ultra2 SCSI 16-Bit High Disk Assembly (# 3023)	411
12.1.10	9.1 GB 10K RPM Ultra2 SCSI Disk Drive Assembly (# 3025)	412
12.1.11	18.2 GB 10K RPM Ultra2 SCSI Disk Assembly (# 3026)	412
12.1.12	9.1 GB 10K RPM Ultra-SCSI Disk Drive (# 3027)	413
12.1.13	9.1 GB Ultra-SCSI 16-bit Enhanced Disk Drive (# 3029)	413
12.1.14	18.2 GB 10K RPM Ultra SCSI Disk Drive (# 3102)	414
12.1.15	18.2 GB Ultra-SCSI Enhanced Disk Drive (# 3103)	414
12.1.16	18.2 GB Ultra-SCSI 1" High Hot-Swap Disk (# 3104)	415
12.1.17	18.2 GB 10K RPM Ultra SCSI Enhanced Drive (# 3110)	415
12.1.18	18.2 GB 10K RPM Ultra SCSI Hot Swap Disk Drive (# 3117)	415
12.1.19	36.4 GB 10K RPM Ultra SCSI Disk (# 3119)	416
12.1.20	36.4 GB 10K RPM Ultra SCSI 16-Bit Enhanced Disk (# 3128)	416
12.1.21	36.4 GB 10K RPM Ultra2 SCSI Disk Assembly (# 3129)	417
12.1.22	9.1 GB SSA 10K RPM Hot Swap Disk (# 3070)	417
12.1.23	18.2 GB SSA 10K RPM Disk Assembly (# 3077)	418
12.1.24	9.1 GB Serial Storage Architecture 10K RPM Disk (# 3079)	418
12.2	CD-ROM Drive	418
12.3	Tape Drives	419
Chapter 13. External Storage Architectures and Devices		421
13.1	Introduction to IBM Open Systems Storage Products	421
13.1.1	IBM Seascope: Storage Enterprise Architecture	421
13.1.2	Storage Area Network	422
13.1.3	RS/6000 and Storage Area Networks	423
13.1.4	Serial Storage Architecture	423
13.2	IBM External Disk Storage	424

13.2.1	Disk Product Positioning	424
13.2.2	Disk Systems Summary	425
13.2.3	IBM 2102 Fibre Channel RAID Storage Server	427
13.2.4	IBM 7133 Serial Disk System Advanced Models T40 and D40	430
13.2.5	Enterprise Storage Servers (2105)	433
13.2.6	Expandable Storage Plus (# 2104)	439
13.2.7	SAN/Fibre Channel Interconnection	440
13.2.8	IBM SAN Fibre Channel Switch and Storage Hub	446
13.2.9	IBM 7204 External Disk Drives	448
13.2.10	IBM 7203 Portable Disk Drive	449
13.3	IBM Tape Drive Products	450
13.3.1	IBM 7205 Digital Linear Tape Drive	451
13.3.2	IBM 7206 External 4 mm DAT Tape Drives	453
13.3.3	IBM 7207 1/4" Cartridge Tape Drives	455
13.3.4	IBM 7208 8 mm Tape Drive	456
13.4	IBM Tape Automation Products	458
13.4.1	Tape Automation Product Positioning Notes	458
13.4.2	Tape Automation Product Summary	458
13.4.3	IBM 3490E Magnetic Tape Subsystem Models	461
13.4.4	IBM Magstar 3590 Tape Subsystem	463
13.4.5	IBM Magstar 3494 Tape Library	466
13.4.6	IBM Magstar MP 3570 Tape Subsystem	471
13.4.7	IBM Magstar MP 3575 Tape Library Dataserver	473
13.4.8	IBM 7332 4 mm DDS Tape Autoloader	475
13.4.9	IBM 7331 8 mm Tape Library	477
13.4.10	IBM 7337 Digital Linear Tape Library	479
13.5	IBM External Optical Devices	480
13.5.1	IBM 7210 CD-ROM Drive	480
13.5.2	IBM 3995 Optical Library Solutions: The Enhanced C-Series	482
13.6	IBM 3466 Network Storage Manager	486
Chapter 14. AIX and the Bonus Pack on the RS/6000		489
14.1	The Base Operating System - AIX	489
14.1.1	Hardware and Software Requirements	489
14.1.2	AIX 4.3 Releases and Support Dates	494
14.1.3	AIX Version 4 Release 3 Features	494
14.1.4	Functions and Utilities - System Management	495
14.1.5	Graphics and GUIs	514
14.1.6	Standards Overview	519
14.1.7	Communications	525
14.1.8	System Performance Enhancements	539
14.1.9	System Security	540
14.1.10	Network Security	545

14.1.11	Performance Analysis Tools Enhancements	546
14.1.12	PC Interoperability - AIX Fast Connect	548
14.1.13	Java Support	550
14.1.14	Additional Printer Support	552
14.1.15	Support and Service	556
14.2	The AIX Bonus Pack - May 7, 2000	562
14.2.1	Database Software	563
14.2.2	Development Tools	564
14.2.3	System Management and Network Management Utilities	565
14.2.4	Multimedia	567
14.2.5	Interoperability Support	567
14.2.6	e-business and Commerce	568
14.2.7	How to Order the Bonus Pack	570
14.3	Third Party Vendors - ISVs	570
Chapter 15. Compilers		571
15.1	C Compiler	571
15.1.1	IBM C for AIX, Version 5 - 5765-E32	571
15.1.2	IBM C for AIX, Version 4.4 - 5765-C64	574
15.1.3	VisualAge C++ Professional, Version 5 - 5765-E26	574
15.1.4	VisualAge C++ Professional for AIX, Version 4 - 5765-D52	576
15.2	Fortran	578
15.2.1	IBM XL Fortran for AIX, Version 7.1	578
15.2.2	XL High Performance Fortran Version 6.1 - 5765-612	580
15.2.3	XL Fortran for AIX	581
Chapter 16. High Availability		583
16.1	High Availability Cluster Multiprocessing	584
16.1.1	Features in HACMP	588
16.1.2	HACMP and HACMP ES Version 4.3	589
16.1.3	HACMP and HACMP ES Version 4.3.1	590
16.1.4	HACMP Version 4.4	592
16.1.5	Upgrade and Migration Paths	593
16.1.6	High Availability Geographic Cluster Product	597
16.1.7	Summary	602
Chapter 17. Parallel Software		605
17.1	Parallel System Support Programs for AIX (PSSP)	605
17.1.1	PSSP Features and Benefits	606
17.1.2	PSSP Prerequisites	607
17.1.3	RS/6000 Cluster Technology	608
17.1.4	Subsystem Communications Support	608
17.1.5	Enabling Parallel Processing	608
17.1.6	Application Programming Models	608

17.1.7	Collecting Performance Data	609
17.1.8	Other Available Services Included	609
17.2	General Parallel File System (GPFS)	609
17.2.1	Administration	610
17.2.2	Standards Compliance	610
17.2.3	Higher Performance and Scalability	610
17.2.4	Availability and Recoverability	610
17.3	LoadLeveler	611
17.3.1	LoadLeveler Highlights	611
17.3.2	LoadLeveler Interfaces	611
17.3.3	What to Expect from LoadLeveler	612
17.4	Parallel Environment	612
17.4.1	Key Elements of PE	613
17.4.2	PE Highlights	613
17.4.3	Parallel Environment Functions	614
17.4.4	Software Requirements	614
17.5	Parallel Engineering and Scientific Subroutine Library (ESSL)	614
17.5.1	Parallel ESSL Highlights	615
17.5.2	Wide Range of Mathematical Functions	615
17.5.3	Application Environments	616
17.5.4	Compatibility with Other Subroutine Libraries	616
17.5.5	SMP Support	616
17.5.6	Useful URLs	617
17.6	Parallel Optimization Subroutine Library (OSL)	617
17.6.1	Parallel Application Programming	618
17.6.2	Linear Programming (LP)	618
17.6.3	Mixed-Integer Programming (MIP)	618
17.6.4	Quadratic Programming (QP)	619
17.6.5	Other Solution Capabilities	619
17.6.6	Parallel OSL Performance	619
17.6.7	Parallel OSL Prerequisites	619
Chapter 18. CATIA Solutions		621
18.1	CATIA Solutions Overview	621
18.2	Operating Environment for CATIA	622
Appendix A. HACMP History		625
A.1	HACMP Version 1.1	625
A.1.1	HACMP Version 1.2	625
A.2	HACMP Version 2.1	626
A.3	HACMP Version 3.1	626
A.3.1	HACMP Version 3.1.1	627
A.4	HACMP Version 4.1	627

A.4.1	HACMP Version 4.1.1	627
A.4.2	HACMP Version 4.2.0	628
A.4.3	HACMP Version 4.2.1	628
A.4.4	HACMP ES Version 4.2.1	629
A.4.5	HACMP and HACMP ES Version 4.2.2	630
A.4.6	HACMP and HACMP ES Version 4.3	630
A.4.7	HACMP and HACMP ES Version 4.3.1	631
A.4.8	HACMP and HACMP ES Version 4.4	632
Appendix B. SCSI Cabling Examples 633		
B.1	Cabling the PCI Single-Ended Ultra SCSI Adapter	633
B.1.1	Cables and Terminators for Single-Ended SCSI Adapters	634
B.2	Cabling the PCI Differential-Ended Ultra SCSI Adapter	637
B.2.1	SCSI Differential Cable Lengths Using This Adapter	637
B.2.2	Cabling Examples for the PCI Differential Ultra SCSI Adapter	640
Appendix C. RS/6000 Performance Data 643		
C.1	SPEC and ROLTP Performance Data - Single Processor Comparison	643
C.2	SPEC and ROLTP Performance Data for SMP Systems	647
C.3	SPECweb96 and SPECweb99 Benchmarks	650
C.4	TPC-C Benchmark Data	651
C.5	TPC-D Published Results (Obsolete as of 8/15/99)	652
C.6	LINPACK Benchmark Data	652
C.7	Notesbench Published Results	653
C.8	Explanations of Performance Benchmarks	653
C.8.1	SPEC Benchmark Definition	654
C.8.2	Relative Online Transaction Processing (ROLTP) Definition	657
C.8.3	Transaction Processing Council (TPC) Benchmark Definitions	658
C.8.4	LINPACK Benchmark Definition	659
C.8.5	Notesbench	659
Appendix D. Adapter Placement Guide 661		
D.1	32-Bit Versus 64-Bit PCI Slots	661
D.2	33 MHz Versus 50 MHz 64-Bit PCI Slots	662
D.3	Connectivity Versus Performance Overview	662
D.4	Other Restrictions	663
D.5	7043 Model 140 Adapter Placement Guide	664
D.6	7043 Model 150 Adapter Placement Guide	667
D.7	7044 Model 170 Placement Guide	671
D.8	7043 Model 260 Adapter Placement Guide	674
D.9	7043/7044 Model 270 Adapter Placement Guide	677
D.10	7046 Model B50 Adapter Placement Guide	680
D.11	7025 Model F50 Adapter Placement Guide	682
D.12	7025 Model F80 Adapter Placement Guide	685

D.13	7026 Model H50 Adapter Placement Guide	687
D.14	7026 Model H70 Adapter Placement Guide	690
D.15	7026 Model H80 Adapter Placement Guide	693
D.16	7026 Model M80 Adapter Placement Guide	696
D.17	7017 Model S70, S70 Advanced, and S80 Adapter Placement Guide	700
D.17.1	Adapter Placement	700
D.17.2	Performance Limits	706
D.17.3	Model S70, Model S70A, and S80 Configuration Details	707
D.18	9076 RS/6000 SP Systems	708
D.19	Digital Trunk PCI Adapter Placement Considerations	709
 Appendix E. Site and Hardware Planning Information		711
E.1	7043 43P Model 140 Site and Hardware Planning Information	711
E.2	7043 43P Model 150 Site and Hardware Planning Information	713
E.3	7044 44P Model 170 Site and Hardware Planning Information	714
E.4	7043 43P Model 260 Site and Hardware Planning Information	716
E.5	7044 44P Model 270 Site and Hardware Planning Information	718
E.6	7046 Model B50 Site and Hardware Planning Information	719
E.7	7025 Model F50 Site and Hardware Planning Information	720
E.8	7025 Model F80 Site and Hardware Planning Information	721
E.9	7026 Model H70 Site and Hardware Planning Information	723
E.10	7026 Model H80 CEC Drawer Site and Hardware Planning Info	724
E.11	7026 Model M80 CEC Drawer Site and Hardware Planning Info	727
E.12	7017 Models S70 and S70 Advanced Site and Hardware Planning	729
E.13	7017 Model S80 Site and Hardware Planning Information	734
E.14	7014 Model S00 Rack Site and Hardware Planning Information	739
E.15	7015 Model R00 Rack Site and Hardware Planning Information	740
E.16	Model T00 and T42 Rack Site and Hardware Planning Information	741
E.17	Noise Emission Notes	742
 Appendix F. Power Cords		743
 Appendix G. Supported System Options		747
 Appendix H. IBM InfoWindow II 3153		751
 Appendix I. Customer Installation Matrix and Processor Groups		753
 Appendix J. The RS/6000 64-Bit Solution		755
J.1	What is 64-Bit Computing?	755
J.1.1	Large File Support	756
J.1.2	Compatibility Starts at the Hardware Level	758
J.1.3	The RS/6000 Enterprise Server Model S70	759
J.1.4	AIX Version 4.3 - Single Product Strategy	760

Appendix K. Special Notices	773
Appendix L. Related Publications	777
L.1 IBM Redbooks	777
L.2 IBM Redbooks Collections	777
L.3 Other Resources	777
L.3.1 RS/6000 Business Partner Sales Kit CD	778
L.4 Referenced Web Sites	779
How to Get IBM Redbooks	781
IBM Redbooks Fax Order Form	782
Abbreviations and Acronyms	783
Index	789
IBM Redbooks Review	813

Figures

1. RS/6000 Machine Time-Line	2
2. Secondary Buses	4
3. Multiple Primary PCI Buses	4
4. General RS/6000 System Block Diagram - Bus-Based Systems	5
5. Advanced RS/6000 System Block Diagram - Switch-Based Systems	6
6. POWER3 Microprocessor Logical Block Diagram	8
7. RS64 Microprocessor Logical Block Diagram	10
8. RS64 III Processor Card	12
9. RS/6000 Systems Handbook Products in Focus	17
10. Model 7043-140 - Front View	28
11. Model 7043-150 with Peripherals	35
12. Model 7044-170 with Peripherals	40
13. Model 7043-260 with Peripherals and Spaceball	47
14. Model 7044-270 with Peripherals	53
15. Model 7046-B50	61
16. Model 7025 F50 - Front View	74
17. Model 7025 F50 - Rear View	76
18. Model 7025 F50 - Internal Disk and Media Bays	83
19. Model 7025 F80 - Front View	85
20. Model 7025 F80 - Rear View	87
21. Model 7025 F80 - Internal Disk and Media Bays	96
22. Model 7026-H50 - Front View in S00 Rack	101
23. Model 7026-H50 - Rear View	103
24. Model 7026-H50 - Internal Media Bays and Disks	104
25. Model 7026-H70 - Front View	110
26. Model 7026-H70 - Rear View	112
27. Model 7026 H70 - Internal Disk and Media Bays	118
28. Model 7026 H80 with I/O drawer - Front View in a Rack	122
29. I/O Drawer - Inner View	125
30. I/O Drawer - Rear View	126
31. RS/6000 Enterprise Server Model M80	137
32. CEC and I/O Rack of the RS/6000 S70 Advanced	157
33. RS64 II Processor Card	160
34. Model S70 Advanced System Switch Complex	161
35. 10 EIA Drawer - Front View	163
36. I/O Drawer Rear View (Partial) with 14 slots and 4 PCI Buses	164
37. I/O Rack and CEC of the Model S80	174
38. RS64 III Processor Card	176
39. Front View of Model S80 CEC	178
40. Rear View of Model S80 CEC	179

41. Model S80 System Switch Complex	181
42. Base SPCN Loop with One I/O Drawer	183
43. Fully Configured SPCN Loop with Four I/O Drawers	183
44. RIO Loop with One I/O Drawer	185
45. One RIO Loop Attached to Two I/O Drawers.	186
46. Two RIO Loops Attached to Two I/O Drawers.	187
47. Two RIO Loops Attached to Four I/O Drawers	188
48. Front View of 10 EIA Drawer	197
49. Rear View of 10 EIA I/O Drawer	198
50. Rack Options Dialog Box.	204
51. Two Systems Sharing One I/O Rack	217
52. Basic Configuration with H50, H70, or H80 Sharing the S80 I/O Rack	218
53. System with One I/O Drawer, Two H50s, H70s, or H80s	219
54. System with Two I/O Drawers with H50, H70, or H80	220
55. Primary and Secondary I/O Drawers in Model S80 I/O Rack	221
56. Model 7014 S00 Rack-Front (Door Removed) and Rear.	227
57. Model 7014-T00 Front View.	230
58. The RS/6000 SP Model 550	243
59. Internal SP Processor Nodes.	253
60. SP Switch Link and SP Switch Port.	287
61. SP Switch Chip	288
62. SP Switch Board	289
63. SP Switch Adapter.	290
64. SP Switch System	291
65. Control Workstation Connectivity and Function	294
66. RS/6000 SP Switch Router Offers Many Possibilities	299
67. SCSI III Standards Overview	394
68. SCSI Signal Comparison	399
69. Piggy-Back Connectors	401
70. Comparison between SCSI and SSA Subsystems	404
71. IBM 2102 Fibre Channel RAID Storage Server	427
72. IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer	430
73. IBM Expandable Storage Plus SCSI	439
74. IBM SAN Data Gateway	441
75. IBM SAN Data Gateway Model S20	444
76. IBM 7205 External Disk Drive	448
77. IBM 7205 Digital Linear Tape Drive.	451
78. IBM 7206 External 4 mm DAT Tape Drive.	453
79. IBM 7207 1/4" Cartridge Tape Drive	455
80. IBM 7208 8 mm Tape Drive.	456
81. IBM 3490E Magnetic Tape Subsystem	461
82. IBM Magstar 3590 Tape Subsystem	463
83. IBM Magstar 3494 Tape Library	466

84. IBM Magstar MP 3570 Tape Subsystem	471
85. IBM Magstar 3575 Tape Library Dataserver	473
86. IBM 7332 Model 220 4 mm Tape Autoloader	475
87. IBM 7331 8 mm Tape Library	477
88. IBM 7337 Digital Linear Tape Library	479
89. IBM 7210 CD-ROM Drive	481
90. IBM 3995 Optical Libraries.	482
91. IBM 3466 Network Storage Manager.	486
92. Costs of Downtime.	583
93. Typical HACMP Solution	586
94. HAGEO Solution	598
95. Differential External Narrow Bus	641
96. Differential External Wide Bus	641
97. Model 140 System Unit Rear View with Numbered Slots	664
98. Model 150 System Unit Rear View with Numbered Slots	667
99. Model 170 System Unit Rear View with Numbered Slots	671
100. Model 260 System Unit Rear View with Numbered Slots	674
101. Model 270 System Unit Rear View with Numbered Slots	678
102. Model B50 System Unit Rear View with Numbered Slots	681
103. Model F50 System Unit Rear View with Numbered Slots	682
104. Model F80 System Unit Rear View with Numbered Slots	685
105. Model H50 System Unit Rear View with Numbered Slots	687
106. Model H70 System Unit Rear View with Numbered Slots	690
107. Model H80 System Unit Rear View with Numbered Slots	693
108. Model M80 System Unit Rear View with Numbered Slots	697
109. I/O Drawer Rear View with 14 Slots and 4 PCI Buses	700
110. 7043 43P Model 140 Physical Planning Diagram	713
111. 7043 43P Model 150 Physical Planning Diagram	714
112. 7044 44P Model 170 Physical Planning Diagram	716
113. 7043 43P Model 260 Physical Planning Diagram	717
114. 7044 44P Model 270 Physical Planning Diagram	719
115. 7025 F50 Physical Planning Diagram	721
116. 7025 F80 Physical Planning Diagram	723
117. 7017 Models S70 and S7A Rack Physical Planning Diagram	733
118. 7017 Models S70 and S7A Rack Physical Planning Diagram	738
119. 7014 Model S00 Rack Physical Planning Diagram	739
120. 7015 Model R00 Physical Planning Diagram	740

Tables

1. Facts and Features for Models 150, 170, and 270	18
2. Facts and Features for Models B50, F50, and F80	19
3. Facts and Features for Models H70 and H80	20
4. Facts and Features for Models M80 and S80	21
5. Facts and Features for SP POWER3 nodes	22
6. Model 140 Standard Configuration	28
7. Model 140 System Expansion	29
8. Model 140 Optional Features	30
9. Publications for the Model 140	34
10. Model 150 Minimum Configuration	35
11. Model 150 System Expansion	36
12. Model 150 Optional Features	37
13. Publications for the Model 150	40
14. Model 170 Minimum Configuration	41
15. Model 170 System Expansion	41
16. Model 170 Optional Features	42
17. Publications for the Model 170	46
18. Additional Publications for the Model 170	46
19. Model 260 Minimum Configuration	47
20. Model 260 System Expansion	48
21. Model 260 Optional Features	49
22. Publications for the Model 260	52
23. Additional Publications for the Model 260	52
24. Model 270 Minimum Configuration	53
25. Model 270 System Expansion	54
26. Model 270 Optional Features	55
27. Publications for the Model 270	60
28. Additional Publications for the Model 270	60
29. Model B50 Minimum Configuration	61
30. Model B50 System Expansion	62
31. Model B50 Optional Features	63
32. Publications for the Model B50	66
33. Additional Publications for the Model B50	66
34. Model F50 Standard Configuration	74
35. Model F50 System Expansion	77
36. Model F50 Optional Features	78
37. Publications for the Model F50	84
38. Model F80 Minimum Configuration	85
39. Model F80 Optional Features	90
40. Publications for the Model F80	99

41. Additional Publications for the Model F80	99
42. Model H50 Standard Configuration	101
43. Model H50 System Expansion	103
44. Model H50 Optional Features	105
45. Publications for the Model H50	109
46. Model H70 Minimum Configuration	111
47. Model H70 System Expansion	112
48. Model H70 Optional Features	113
49. Model H80 Minimum Configuration	123
50. Model H80 Optional Features	127
51. Publications for the Model H80	134
52. Additional Publications for the Model H80	134
53. Model M80 Standard Configuration	141
54. Model M80 System Expansion	141
55. Model M80 Optional Features	142
56. Publications Shipped with the Model M80	150
57. Model S70 Advanced Minimum Configuration	166
58. Model S70 Advanced System Expansion	167
59. Model S70 Advanced Optional Features	168
60. Publications for the Model S70 Advanced	172
61. Model S80 Minimum Configuration	190
62. Model S80 System Expansion	191
63. Model S80 Optional Features	191
64. S70 Advanced to S80 Processor Conversion	201
65. S70 Advanced to S80 Memory Conversion	201
66. Publications Shipped with the Model S80	205
67. HA-S7A Minimum Configuration	209
68. Model 7015-R00 Net Weight and Physical Size	224
69. Publications Shipped with the Model R00	226
70. Model 7014-S00 Net Weight and Physical Size	227
71. Publications Available for the Model 7014-S00	229
72. Model 7014-T00 Net Weight and Physical Size	233
73. Publications Available for the Model 7014-T00	235
74. Model 7014-T42 Net Weight and Physical Size	239
75. Publications Available for the Model 7014-T42	241
76. SP Legacy Models by Frame and Switch Types	248
77. SP Legacy Models by Node Types	249
78. SP Frame Dimensions	250
79. SP Frame Electrical Power Requirements	251
80. Standard Configurations for POWER3-II 375 MHz SMP Nodes	255
81. POWER3 SMP 222 MHz High Node	256
82. Standard Configuration for POWER3 200 MHz SMP Nodes	257
83. Standard Configuration for the 332 MHz PowerPC 604e Nodes	258

84. POWER3 SMP 222 MHz High Node Memory Options	261
85. Internal Disk on POWER3 222 MHz High Node or its Expansion I/O	262
86. Hot Plug PCI Adapters for 375 MHz POWER3-II Wide Node	264
87. Minimum and Maximum Memory	268
88. Potential Disk pairs in 375 MHz Wide Node	269
89. Processor Options for 332 MHz SMP Wide Nodes	273
90. Memory Features for 332 MHz SMP Wide Nodes	274
91. Minimum and Maximum Memory	279
92. Potential Disk Pairs in 375 MHz Thin Node	279
93. Processor Options for 332 MHz SMP Thin Nodes	283
94. Memory Features for 332 MHz SMP Thin Nodes	284
95. Supported and Unsupported SP Control Workstations	295
96. Supported Communications PCI Adapters on the SMP SP Nodes	307
97. AIX and PSSP Levels Supported on SP Nodes	309
98. SP Nodes Not Current	310
99. Specification Summary of IBM NUMA-Q 2000 Model E410	318
100. IBM NUMA-Q Cabinet	320
101. NUMA-Q Server Summary	320
102. Available PCI I/O Adapters and their Feature Codes	323
103. The 1.2 Mbps RANs and Cables	334
104. The 2.4 Mbps RANs and Cables	334
105. ARTIC960Hx Cables	336
106. Maximum Port Speed for Each Electrical Interface	337
107. ARTIC960 Cables and Features	337
108. Cable Information for the 2-Port Multiprotocol PCI Adapter	339
109. Overview of the Compatibility and Use of Both Optical Extenders	362
110. Hot Pluggable PCI Adapters	364
111. GXT3000P Display and Cable Table	371
112. GXT3000P Performance Table - Historic Measures	372
113. GXT3000P Performance Table - Current Measures	372
114. GXT2000P Display and Cable Table	374
115. GXT2000P Performance Table - Historic Measures	375
116. GXT2000P Performance Table - Current Measures	375
117. GXT300P Display and Cable Table	377
118. GXT300P Performance Table	378
119. GXT250P and GXT255P Display and Cable Table	380
120. GXT130P Display and Cable Table	382
121. GXT130P Performance Table	383
122. GXT120P Display and Cable Table	385
123. Performance Improvement in Percentages	386
124. Non-Supported Graphics Accelerators	387
125. Maximal Number of Accelerators Per System and Limitations	389
126. Overview of SCSI-III Standards	394

127.SCSI Specification Summary	400
128.Comparing SSA and SCSI	404
129.Available Disk Drives and Feature Codes	407
130.Internal 32X CD-ROM Drive	419
131.Internal 4 mm and 8 mm Tape Drives	419
132.IBM Open Systems Storage External Disk Systems	425
133.2102 IBM Fibre Channel RAID Storage Server	429
134.IBM 7133 Serial Disk System Advanced Models at a Glance	432
135.Model Summary Matrix	434
136.List of RS/6000 Systems	436
137.ESS Characteristics at a Glance	437
138.HACMP hardware support for ESS models	438
139.Available Products Matrix for the IBM Expandable Storage Plus	440
140.IBM SAN Data Gateway for Serial Disk at a Glance	446
141.IBM 7204 External Disk Drives at a Glance	448
142.IBM Tape Drive Products	450
143.IBM 7205-311 Digital Linear Tape Drive at a Glance	451
144.IBM 7206-110 DDS-3 DAT Tape Drive at a Glance	453
145.IBM 7206 Model DDS-4 at a Glance	454
146.IBM 7207-122 4 GB SLR5 Tape Drive at a Glance	455
147.IBM 7208-341 8 mm Tape Drive at a Glance	456
148.IBM Tape Automation Products	459
149.IBM Tape Automation (Continued)	460
150.IBM 3490E Model F01/F11 Magnetic Tape Subsystem at a Glance	462
151.IBM Magstar 3494 Tape Library at a Glance	469
152.IBM Magstar MP 3570 Tape Subsystem at a Glance	472
153.IBM Magstar MP 3575 Tape Library Dataserver at a Glance	474
154.IBM 7331 Model 305 8 mm Tape Library at a Glance	478
155.IBM 7337 Digital Linear Tape Library at a Glance	480
156.IBM 7210 Model 015 External CD-ROM Drive at a Glance	482
157.IBM 3995 Optical SCSI Models at a Glance	483
158.IBM 3466 Network Storage Manager at a Glance	487
159.Memory and Disk Requirements	492
160.AIX 4.3 Version Availability Dates	494
161.Differences between NIS and NIS+	526
162.AIX Publications	559
163.Bonus Pack Ordering Requirements	570
164.C Versions and Upgrade Paths	577
165.Fortran Versions and Their AIX Levels	581
166.Upgrade Path or Migration for HACMP Classic	594
167.Compatibility of HACMP with Supported AIX Version 4	594
168.Upgrade or Migration for HACMP ES	595
169.Compatibility of HACMP ES with AIX Version 4	595

170.Migration from HACMP Classic to HACMP ES	596
171.Supported Systems, Devices, or Features with HACMP Versions	602
172.Supported Systems, Devices, or Features with HACMP Versions	603
173.Single-Ended Ultra SCSI Adapter-to-First Device Cables.	635
174.Device-to-Device Cables for Single-Ended Installations	636
175.16-Bit SCSI-II System-to-System Cable	637
176.Terminators for Single-Ended Installations	637
177.Differential-Ended Ultra SCSI Adapter-to-First Device Cables	638
178.Device-to-Device Cables for Differential-Ended Installations	638
179.Terminator for Description-Ended Installations	640
180.Cabling for the PCI Differential Ultra SCSI Adapter	641
181.SPEC and ROLTP Performance Data for Uniprocessor Systems	643
182.SPEC2000 for Uniprocessor Systems	646
183.SPEC and ROLTP Performance Data for SMP Systems	647
184.SPEC2000 for SMP Systems	650
185.Web Benchmarks for Selected RS/6000 Systems	650
186.TPC-C Benchmark Results	651
187.TPC-D Published Results	652
188.TPC-D Published Results for SP Nodes	652
189.LINPACK Performance Data for Uniprocessor Systems.	652
190.Notesbench Published Results	653
191.Model 140 Single Adapter Placement Guide	664
192.Model 140 Multiple Adapter Placement Guide	666
193.Model 150 Single Adapter Placement Guide	668
194.Model 150 Multiple Adapter Placement Guide	670
195.Model 170 Single Adapter Placement Guide	672
196.Model 170 Multiple Adapter Placement Guide	673
197.Model 260 Single Adapter Placement Guide	675
198.Model 260 Multiple Adapter Placement Guide	676
199.Model 270 Single Adapter Placement Guide	678
200.Model 270 Multiple Adapter Placement Guide	680
201.Communication, Graphics, and Storage Adapters Summary	681
202.Model F50 Adapter Placement Guide	683
203.Model F80 Adapter Placement Guide	685
204.Model H50 Adapter Placement Guide.	688
205.Model H70 Adapter Placement Guide.	690
206.Model H80 Adapter Placement Guide.	694
207.Model M80 Adapter Placement Guide	697
208.Models S70,S70 Advanced Adapter, and S80 Placement Guide	701
209.S70, S7A, and S80 Performance Limits	707
210.7043 43P Model 140 Site and Hardware Planning Information	711
211.7043 43P Model 150 Site and Hardware Planning Information	713
212.7044 44P Model 170 Site and Hardware Planning Information	714

213.7043 43P Model 260 Site and Hardware Planning Information	716
214.7044 44P Model 270 Site and Hardware Planning Information	718
215.7046 Model B50 Site and Hardware Planning Information	719
216.7025 Model F50 Site and Hardware Planning Information	720
217.7026 Model F80 Site and Hardware Planning Information	721
218.7026 Model H70 Site and Hardware Planning Information	723
219.7026 Model H80 CEC Drawer Site and Hardware Planning Info	724
220.Site and Hardware Planning Information for I/O Drawer	726
221.7026 Model M80 CEC Drawer Site and Hardware Planning Info	727
222.Site and Hardware Planning Information for I/O Drawer	728
223.7017 Models S70 and S7A Rack Physical Planning Information	730
224.7017 Model S70 SCSI I/O Drawer 7 EIA Planning Information.	731
225.Model S70 Advanced I/O Drawer 10 EIA Planning Information	732
226.7017 Models S70 and S70 Advanced I/O Rack Planning Information . .	733
227.7017 Model S80 CEC Rack Physical Planning Information	735
228.7017 Model S80 SCSI I/O Drawer 7 EIA Planning Information.	736
229.Model S70 Advanced I/O Drawer 10 EIA Planning Information	737
230.7017 Models S70 and S70 Advanced I/O Rack Planning Information . .	738
231.7014 Model S00 Rack Site and Hardware Planning Information	739
232.7015 Model R00 Rack Site and Hardware Planning Information	740
233.Model T00 Rack Site and Hardware Planning Information	741
234.Model T42 Rack Site and Hardware Planning Information	742
235.Power Cord Connector Spotters Reference	743
236.External Devices for RS/6000 Models 140, 150, 170, 260, 270, and B50747	
237.External Devices for RS/6000 Models F50, F80, H50, H70, and M80 . .	748
238.External Devices for RS/6000 Models S7A, S80, and SP.	749
239.RS/6000 Models and Feature Codes Set Up Information.	753
240.AIX Programming Environments	762
241.AIX Virtual Address Space Layout	767

Preface

The *RS/6000 Systems Handbook 2000 Edition*, SG24-5120, is a comprehensive single-source guide covering the RS/6000 product line. Major hardware and software offerings are introduced and their prominent functions discussed.

This publication is suitable for professionals wishing to acquire a better understanding of RS/6000 products, including:

- Customers
- Sales and marketing professionals
- Technical support professionals
- IBM Business Partners

Inside this publication, you will find:

- A historical look at RS/6000 hardware
- An overview of the latest RS/6000 models
- A short discussion on hardware architecture
- Information on storage, graphics, and communications features
- A description of AIX and supported software platforms
- Hundreds of tables and figures providing effective access to useful information
- A special chapter on NUMA-Q to introduce these servers.

The introduction of this redbook expands the current set of RS/6000 handbooks by providing an ideal, comprehensive, desktop reference that covers the entire product range from the desktop to the raised floor.

This publication does not replace the latest RS/6000 marketing materials and tools. It is intended as an additional source of information that, together with existing sources, may be used to enhance your knowledge of IBM's solutions for the UNIX marketplace.

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.

Jana Babnik-Gomiscek is an I/T specialist working in IBM International Technical Center in Ljubljana - Slovenia, and she is responsible for technical support in Central and Eastern Europe/Middle East/Africa (CEMA). She is also Business Partner Technical Manager for CEMA, RS/6000 Team Leader, and a member of the High-End Technology Focus Group. She holds a degree in Computer Science from the University of Ljubljana and a Ph.D. in Parallel Computing from the Technical University of Vienna.

Volker Haug is an Advisory I/T Specialist. He has more than thirteen years of experience in the I/T industry, the last ten of which he has devoted to RS/6000 and AIX systems. He holds a degree in Business Management from the Berufsakademie in Stuttgart. Volker is the technical leader supporting IBM sales, business partners, and customers with pre-sales consultation and implementation of client/server environments. His areas of expertise include RS/6000 workstations and workgroup servers, graphics, MCAD applications, and AIX systems management. Based in Stuttgart, Germany, he is currently working for the Web Server Sales RS/6000 and NUMA-Q technical support organization covering the IBM Europe, Middle East, and Africa Central Region.

Jeanine Indest is an I/T Specialist with the Techline RS/6000 Sales Support Team, part of the Web Server Sales Americas Organization. She has more than six years of experience in the I/T industry, the last four of which she has devoted to RS/6000 and AIX systems. She holds a B.S. degree and a M.A. degree in Economics from the University of New Orleans. Her current area of expertise is technical marketing support for Business Partners in the United States for RS/6000 and NUMA-Q systems.

Stephen Lutz is an IT Specialist in the technical support for RS/6000 and NUMA-Q, part of the Web Server Sales Organization in Stuttgart, Germany. He holds a degree in Computer Science from the Fachhochschule Karlsruhe - University of Technology and is an IBM Certified Advanced Technical Expert. Stephen is a member of the High-End Technology Focus Group, supporting IBM sales, Business Partners, and customers with pre-sales consultation and implementation of client/server environments.

Shyam Manohar is a Marketing Manager from IBM-India with the Web Server Group in ESG. Shyam is one among the leading top contributors from India for the IBM-ASEAN region. With his valuable experience in the RISC/UNIX marketplace, Shyam holds unique expertise in product positioning, especially in competitive scenarios. He is part of a focus team engaged in large, high-end server opportunities.

Irene D. Sideris is an I/T Specialist serving in the RS/6000 Techline in Toronto, Canada. She has eleven years of experience in the RS/6000 field, having worked seven years in the Customer Assist Center and four years in Techline. She holds an honors B.A. from the University of Toronto, and a B.Sc. from Devry Institute of Technology. She is an IBM Certified Specialist in AIX Systems Administration, User, and RS/6000 Solutions Sales.

The authors of the first edition are:

Jennifer Acuna-Narvaez	San Francisco, U.S.A.
Ashoka Reddy	Basingstoke, U.K.
Linda Sandberg	IBM Stockholm, Sweden
Irene D. Sideris	IBM Toronto, Canada

The project that produced this publication was coordinated by:

Scott Vetter	IBM Austin, U.S.A.
---------------------	--------------------

Thanks to the following people for their invaluable contributions to this project:

Larry Amy	IBM Austin, U.S.A.
Ricardo Aranda	IBM Austin, U.S.A.
Jim Beesley	IBM Austin, U.S.A.
Alice Bird	IBM San Jose, U.S.A.
Raman Chhibbar	IBM Toronto, Canada
Teri Dewalt	IBM Austin, U.S.A.
Herman Dierks	IBM Austin, U.S.A.
Funda Eceral	IBM San Jose, U.S.A.
Rich Enders	IBM Austin, U.S.A.
Bill Flint	IBM Austin, U.S.A.
Kay Freund	IBM Austin, U.S.A.
Brad Frey	IBM Austin, U.S.A.
Fred Hudson	IBM Austin, U.S.A.
William Klaiber	IBM Tucson, U.S.A.
Alan MacKay	IBM Austin, U.S.A.
Nikhil A. Mehta	IBM Austin, U.S.A.
Dave Marquardt	IBM Austin, U.S.A.

Bill Mihaltse	IBM Somers, U.S.A.
George L. More	IBM Austin, U.S.A.
Thoi Nguyen	IBM Austin, U.S.A.
Amy O'Brien	IBM Beaverton, U.S.A.
Ron Oreshan	IBM Austin, U.S.A.
Raj Patel	IBM Austin, U.S.A.
Duke Paulsen	IBM Austin, U.S.A.
Bruce Prewit	IBM Austin, U.S.A.
Robert Puckett	IBM Austin, U.S.A.
Doug Ranz	IBM Southfield, U.S.A.
Dusty Rhodes	IBM Austin, U.S.A.
Norm Roller	IBM Austin, U.S.A.
David Ruth	IBM Austin, U.S.A.
Carolyn Scherrer	IBM Austin, U.S.A.
Craig Shempert	IBM Austin, U.S.A.
Nancy Springen	IBM Austin, U.S.A.
Bill Stayer	IBM Beaverton, U.S.A.
Mike Stys	IBM Austin, U.S.A.
Frances Tan	IBM San Jose, U.S.A.
Margot Theriault	IBM Calgary, Canada
Teerasit Tinnakul	IBM Austin, U.S.A.
Tom Torrico	IBM Austin, U.S.A.
Ed Toutant	IBM Austin, U.S.A.
Brian Vicknair	IBM Austin, U.S.A.
Tomaz Vincek	IBM Montpellier, France
Joseph Wilding	IBM Austin, U.S.A.

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 “IBM Redbooks Review” on page 813 to the fax number shown on the form.
- Use the online evaluation form found at ibm.com/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. RS/6000 Introduction

This chapter introduces RS/6000, giving a brief history of the products, an overview of the RS/6000 design, and a description of key RS/6000 technologies.

The RS/6000 family combines the benefits of UNIX computing with IBMs leading-edge RISC technology in a broad product line - from powerful desktop workstations ideal for mechanical design, to workgroup servers for departments and small businesses, to enterprise servers for medium to large companies for ERP and server consolidation applications, up to massively parallel RS/6000 SP systems that can handle demanding scientific and technical computing, business intelligence, and Web serving tasks. Along with AIX, IBMs award winning UNIX operating system, and HACMP, the leading high availability clustering solution, the RS/6000 platform provides the power to create change and has the flexibility to manage it with a wide variety of applications that provide real value.

1.1 RS/6000 History

The first RS/6000 was announced February 1990 and shipped June 1990. Since then, over 1,100,000 systems have shipped to over 132,000 customers.

Figure 1 on page 2 summarizes the history of the RS/6000 product line, classified by machine type. For each machine type, the I/O bus architecture and range of processor clock speeds are indicated. This figure shows the following:

- In the past, RS/6000 I/O buses were based on the Micro Channel Architecture (MCA). Today, RS/6000 I/O buses are based on the industry-standard Peripheral Component Interface (PCI) Architecture.
- Processor speed, one key element of RS/6000 system performance, has increased dramatically over time.
- There have been many machine types over the entire RS/6000 history. In recent years, there has been considerable effort to reduce the complexity of the model offerings without creating gaps in the market coverage.

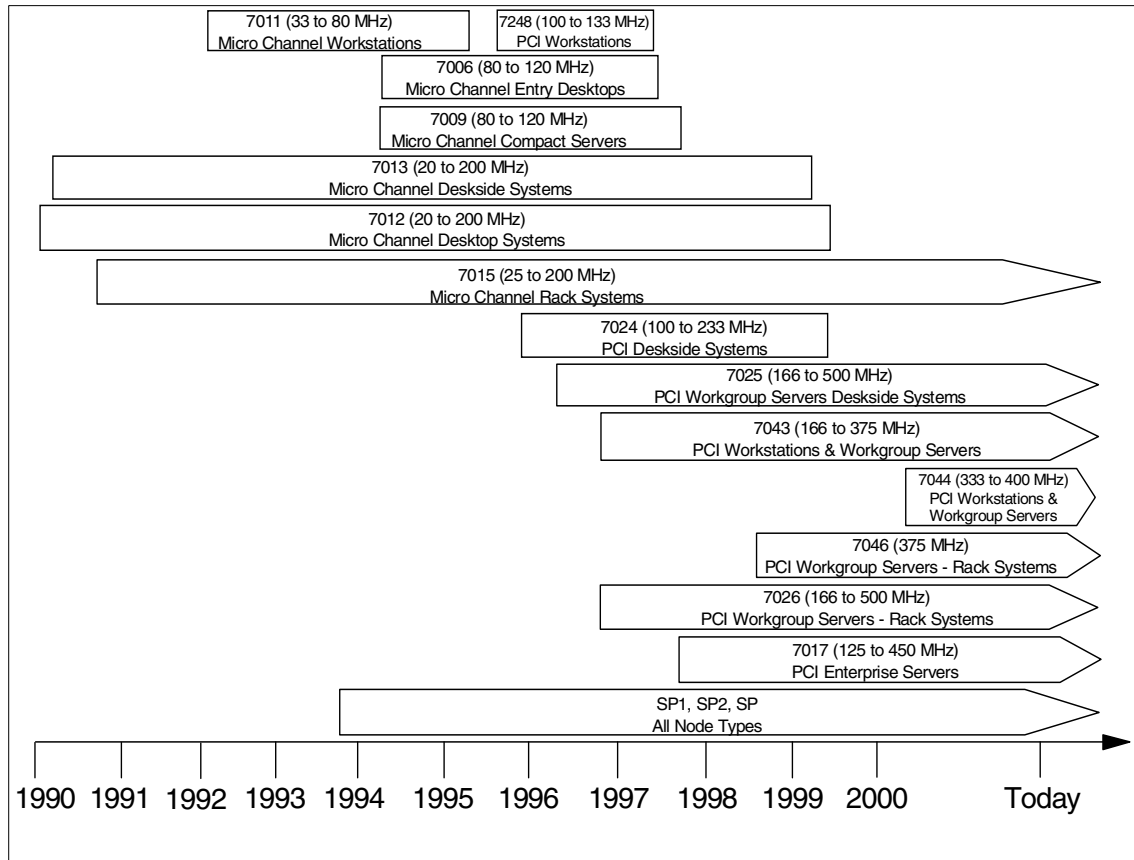


Figure 1. RS/6000 Machine Time-Line

1.2 RS/6000 Design Overview

This section provides information on the following elements, which are important in the design of RS/6000 machines:

- A general system block diagram explanation
- The RS/6000 microprocessor architectures

1.2.1 Computer System Block Diagram Explanation

Figure 4 on page 5 shows a general RS/6000 system block diagram for bus-based systems, such as the Model 270. Figure 5 on page 6 shows the advanced design of RS/6000 switch-based systems, such as the Model M80 or Model S80.

Specific details in the figures, along with the supporting discussions, may not apply to all RS/6000 models covered in this publication. However, the general architecture discussion helps build an overall appreciation for RS/6000 hardware.

All platforms (from workstations to high-end servers) consist of one or more microprocessors, a volatile system memory separate from other subsystems, and a number of I/O devices that may initiate transactions to system memory.

The processors are linked over the primary processor bus/switch to each other, to the system memory, and to one or more host bridges.

In general, I/O devices do not connect to the primary processor bus/switch. The host bridges connect to secondary buses that have I/O devices connected to them. Most commonly, the adapters use the PCI architecture.

1.2.1.1 PCI Slots

The PCI architecture provides an industry standard specification and protocol that allows multiple adapters access to system resources through a set of adapter slots.

Each PCI bus has a limit on the number of slots (adapters) it can support. Typically, this can range from two to six. To overcome this limit, the system design can implement multiple PCI buses. Two different methods can be used to add PCI buses in a system. These two methods are:

- Add secondary PCI buses off the primary PCI bus
- Implement multiple primary buses.

1.2.1.2 Secondary PCI Bus

The simplest method to add PCI slots when designing a system is to add a secondary PCI bus. This bus is *bridged* onto a primary bus using a PCI-to-PCI bridge chip. An example of this design is shown in Figure 2 on page 4.

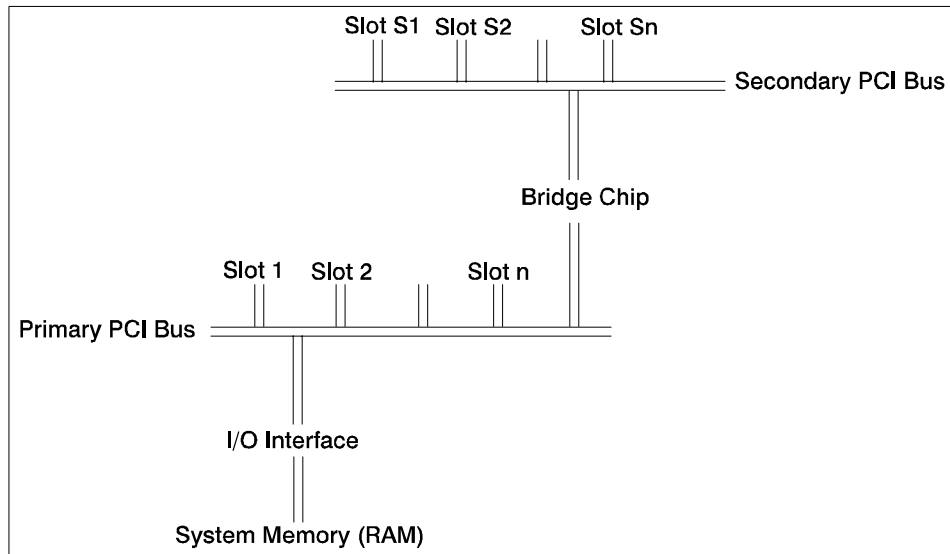


Figure 2. Secondary Buses

Because the slots on the secondary PCI bus must pass through the bridge chip, certain adapters on a secondary PCI bus may experience lower performance. An example of this is the 7025-F40, which was designed with a secondary PCI bus.

1.2.1.3 Multiple Primary PCI Buses

Another method of providing more PCI slots is to design the system with two or more primary PCI buses. This design requires a more sophisticated I/O interface with the system memory. An example of this design is shown in Figure 3. The 7044-270 is designed using multiple primary PCI buses.

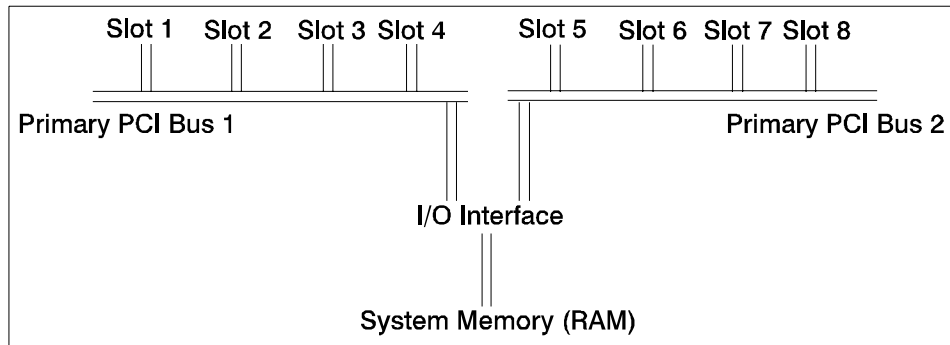


Figure 3. Multiple Primary PCI Buses

This design has improved I/O performance over the secondary bus method because there are multiple parallel paths into the system memory.

The design in Figure 4 shows a peer bus implementation, which allows all the PCI buses to have equal access to the memory controller.

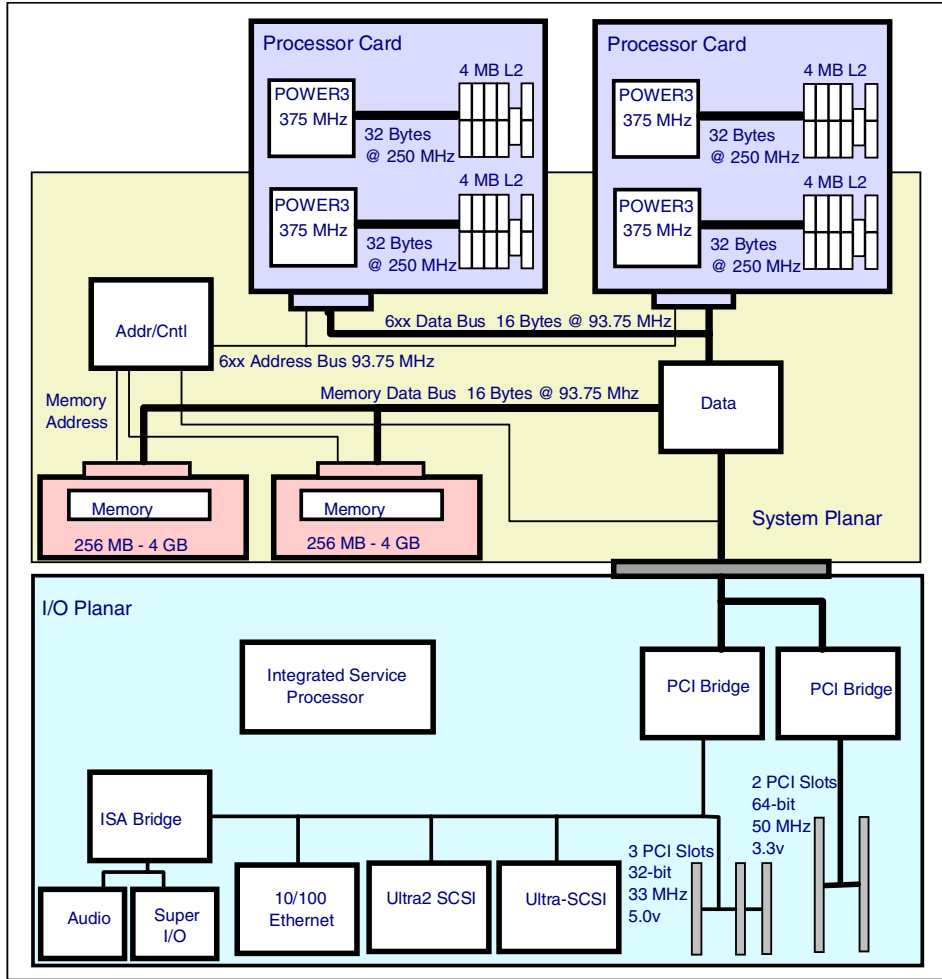


Figure 4. General RS/6000 System Block Diagram - Bus-Based Systems

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

Switch-based memory controllers can provide a peak aggregate throughput of 19.2 GB/s, thus providing superior performance over bus-based systems when the number of processors and I/O connections increase. Figure 5 shows a switch-based design using multiple memory controllers.

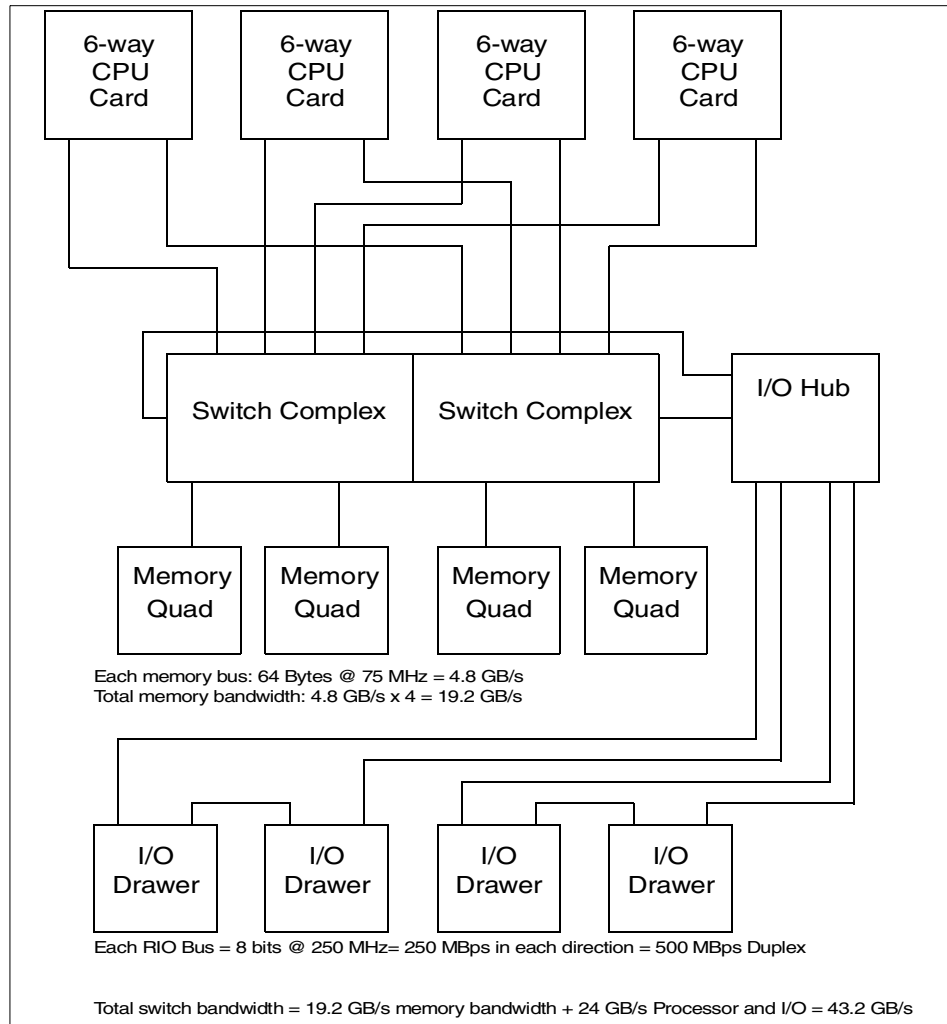


Figure 5. Advanced RS/6000 System Block Diagram - Switch-Based Systems

1.2.1.4 Integrated Adapters

A number of devices are now integrated onto the main processor board, but they physically connect to one of the PCI buses. For this reason, some of the

buses may only have two or three slots available to install adapters. Examples of integrated PCI adapters are: SCSI adapters or Ethernet adapters.

1.2.2 RS/6000 Microprocessor Architectures

Though the entire system architecture contributes to the performance of the RS/6000 product line, the processors are a key component of system performance. The following sections outline the architectures of the latest RS/6000 microprocessors.

1.2.2.1 POWER3 and POWER3-II Microprocessor

The POWER3 microprocessor introduces a new generation of 64-bit processors especially designed for high performance and visual computing applications. POWER3 processors replace the POWER2 and the POWER2 Super Chips (P2SC) in high-end RS/6000 workstations and SP nodes. The RS/6000 44P 7044 Model 270 workstation features the POWER3-II microprocessor as well as the POWER3-II based SP nodes.

The POWER3 implementation of the PowerPC architecture provides significant enhancements compared to the POWER2 architecture. The SMP-capable POWER3 design allows for concurrent operation of fixed-point instructions, load/store instructions, branch instructions, and floating-point instructions. Compared to the P2SC, which reaches its design limits at a clock frequency of 160 MHz, POWER3 is targeting up to 600 MHz by exploiting more advanced chip manufacturing processes, such as copper technology. The first POWER3-based system, RS/6000 43P 7043 Model 260, runs at 200 MHz as well as the POWER3 wide and thin nodes for the SP.

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

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

Figure 6 on page 8 shows the POWER3 microprocessor architecture.

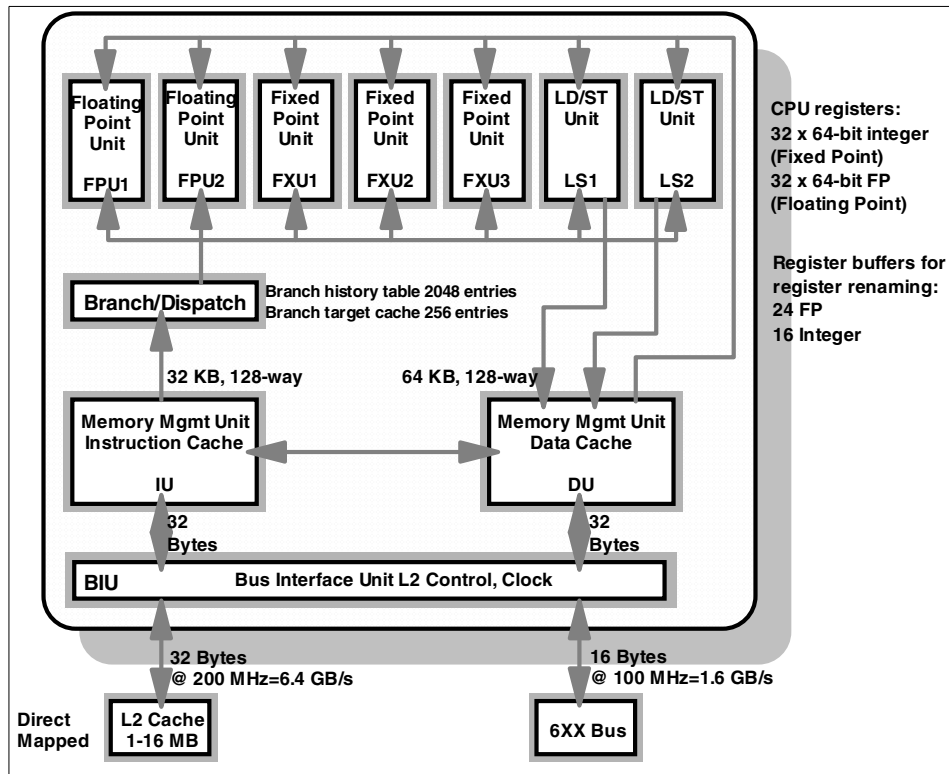


Figure 6. POWER3 Microprocessor Logical Block Diagram

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

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

POWER3 is defined by the following specifications:

- POWER3 microprocessor running at:
 - 200 MHz on RS/6000 Model 260
 - 200 MHz on RS/6000 SP POWER3 SMP Thin Node

- 200 MHz on RS/6000 POWER3 SMP Wide Node
- Superscalar design with integrated integer, floating-point, and branch units
- 32 KB instruction cache
- 64 KB 128-way set associative data cache
- 128-bit memory interface with 64-bit addressing
- Real memory support for up to 4 GB (3 GB on AIX 4.2.1)
- CMOS 6S2 using a 270 mm² die
- 15 million transistors

On February 7, 2000, IBM announced the 44P family and a new SP Node, both of which are based on the POWER3-II technology. POWER3-II is an enhancement to the POWER3 architecture, and it runs with up to 400 MHz and is based on copper technology (CMOS7S). The reduction of the die size made available using CMOS7S technology allowed the addition of an improved L2 data cache manager, additional clock to bus ratios, and the ability to run up to 8-way SMP combinations.

1.2.3 RS64 and RS64 II Microprocessors

The RS64 microprocessor, based on the PowerPC Architecture, was designed for leading-edge performance in OLTP, e-business, BI, server consolidation, SAP, Notesbench, and Web serving for the commercial and server markets. It is the basis for at least four generations of RS/6000 and AS/400 enterprise server offerings.

The RS64 processor block diagram shown in Figure 7 on page 10 focuses on commercial performance with emphasis on conditional branches with zero or one cycle incorrect branch predict penalty, contains 64 KB L1 instruction and data caches, has a one cycle load support, four superscalar fixed point pipelines, and one floating point pipeline. There is an on-board bus interface (BIU) that controls both the 32 MB L2 bus interface and the memory bus interface.

RS64 and RS64 II are defined by the following specifications:

- 125 MHz RS64/262 MHz RS64 II on the RS/6000 Model S70
- 262 MHz RS64 II on the RS/6000 Model S70 Advanced
- 340 MHz RS64 II on the RS/6000 Model H70
- 64 KB on-chip, L1 instruction cache
- 64 KB on-chip four-way set associative data cache

- 32 MB L2 cache
- Superscalar design with integrated integer, floating-point, and branch units
- Support for up to 64-way SMP configurations (currently 12-way)
- 128-bit data bus
- 64-bit real memory addressing
- Real memory support for up to one terabyte (2^{40})
- Virtual memory support for up to one yottabyte (2^{80})
- CMOS 6S2 using a 162 mm² die
- 12.5 million transistors

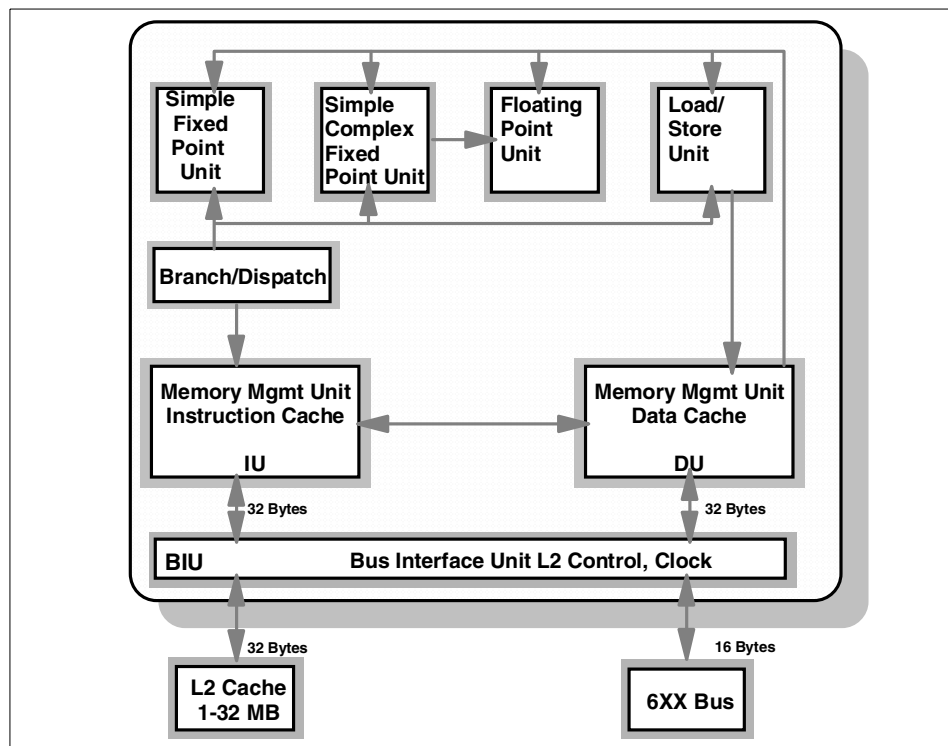


Figure 7. RS64 Microprocessor Logical Block Diagram

1.2.4 RS64 III Processor

The RS64 III processor is designed to perform applications that place heavy demands on system memory. The RS64 III architecture addresses both the need for very large working sets and low latency. Latency is measured by the

number of CPU cycles that elapse before requested data or instructions can be utilized by the processor.

The RS64 III processors combine IBM advanced copper chip technology with a redesign of critical timing paths on the chip to achieve greater throughput. The L1 instruction and data caches have been doubled to 128 KB each. New circuit design techniques were used to maintain the one cycle load-to-use latency for the L1 data cache.

L2 cache performance on the RS64 III processor has been significantly improved. Each processor has an on-chip L2 cache controller and an on-chip directory of L2 cache contents. The cache is four-way set associative. This means that directory information for all four sets is accessed in parallel. Greater associativity results in more cache hits and lower latency, which improves commercial performance.

Using a technique called Double Data Rate (DDR), the new 8 MB Static Random Access Memory (SRAM) used for L2 is capable of transferring data twice during each clock cycle. The L2 interface is 32 bytes wide and runs at 225 MHz (half processor speed), but, because of the use of DDR, it provides 14.4 GBps of throughput.

Figure 8 on page 12 contains a diagram of the processor card layout used on the Model S80. There are six RS64 III processors per CPU card sharing two system buses that connect to the memory controller complex.

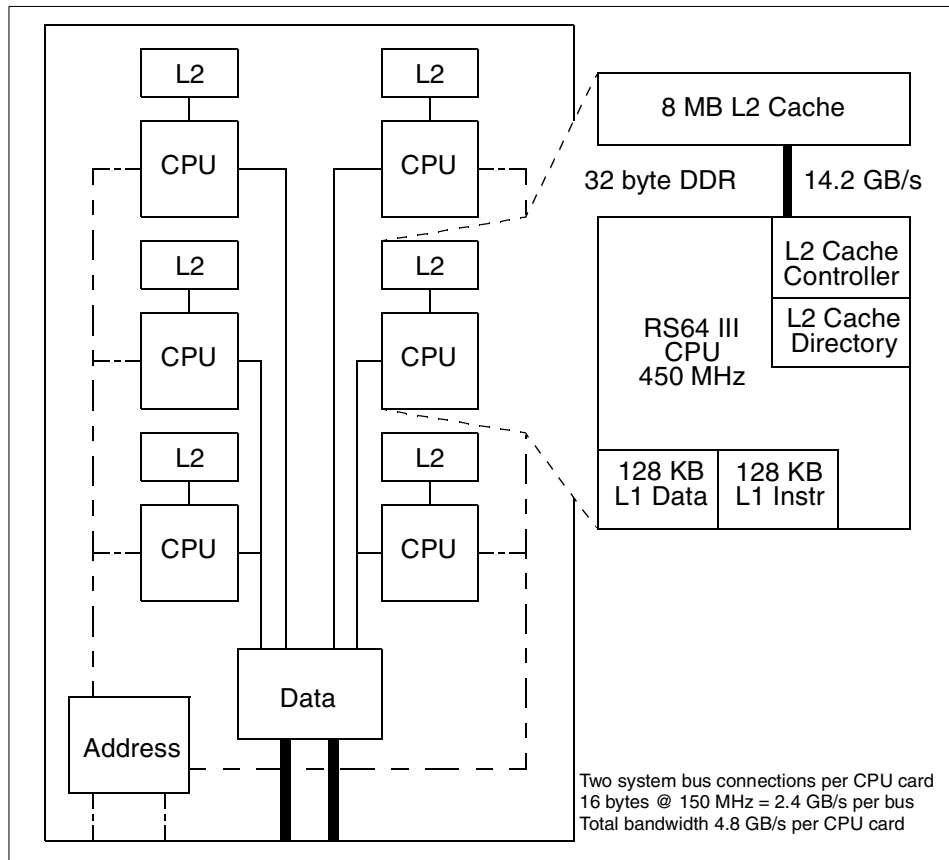


Figure 8. RS64 III Processor Card

The RS64 III processor has five pipeline execution units: Branch, load/store, fixed point, complex fixed point, and floating point. The complex fixed point unit provides support for multiplication and division math. There is a dispatch buffer that can hold up to 16 current instructions, a technique that helps reduce latency. It also has an eight-deep branch buffer. The processor can sustain a decode and execution rate of up to four instructions per cycle.

All processor arrays use fault detection and correction techniques, such as redundancy, error checking and correction (ECC), parity, and retry. Together, these tools promote high reliability, availability, and data integrity. This enables full fault detection and correction coverage within the CEC.

In summary, the RS64 III features include:

- 128 KB on-chip L1 instruction cache

- 128 KB on-chip L1 data cache with one cycle load-to-use latency
- On-chip L2 cache directory that supports up to 8 MB of off-chip L2 SRAM memory
- 14.4 GBps L2 cache bandwidth
- 32 byte on-chip data buses
- 4-way superscalar design
- Five stage deep pipeline
- The Model S80 uses the 450 MHz RS64 III 64-bit copper-chip technology
- The Model M80 uses the 500 MHz RS64 III 64-bit copper-chip technology
- The Model F80 and the H80 use 450 or 500MHz RS64 III 64-bit copper-chip technology

1.3 RS/6000 New Technologies and Directions

This section gives a brief overview of the following key, leading-edge RS/6000 technologies and directions:

- 64-bit technology
- Processor technologies, namely copper circuitry and silicon-on-insulator (SOI)
- IBM POWER4 microprocessor
- AIX on Intel support (Project Monterey)

1.3.1 64-Bit Technology

64-bit computing is the direction for all RS/6000 products. The essence of the RS/6000 64-bit computing strategy can be summed up in three themes:

1. Complementing the established scalability of the existing 32-bit product set, 64-bit technology is the enabler for scaling enterprise SMP servers to higher capacity, making high-end system performance one of the primary customer benefits of 64-bit computing. This enabling may be summarized as: 64-bit hardware expands addressability; addressability creates the potential for higher capacity; and a balanced 64-bit system design will realize the potential by delivering higher system performance.
2. 64-bit computing is complementary to 32-bit computing. Customers want the benefits of 64-bit technology available to them but know their 32-bit systems and 32-bit applications will be important investments for a long time. A strategic focus on the complementary nature of the new

technology forms the thread connecting many RS/6000 and AIX Version 4.3 64-bit features.

3. The transition from 32-bit computing to a future in which 64-bit and 32-bit computing coexist will be, for RS/6000 customers, a very smooth evolution.

These themes were also the design principles behind the product implementation of AIX Version 4.3 as a 64-bit operating environment. AIX Version 4.3 introduces significant functional and scalability enhancements that benefit all RS/6000 customers. A single AIX product supporting both 64-bit computing, as well as broad general improvements, is a prime example of the RS/6000 evolutionary vision at work.

1.3.2 Processor Technologies

IBM has developed industry-leading microprocessor fabrication technologies. These technologies are copper circuitry and silicon-on-insulator (SOI) on complimentary metal oxide semiconductor (CMOS) chips.

As an electrical conductor, copper is superior to aluminum, which has been used for the past 30 years. The net effect of using copper circuitry is increased clock speeds, smaller die sizes, smaller channel lengths, and lower voltages.

SOI protects the millions of transistors on a chip with a thin layer of silicon oxide, thus reducing harmful electrical effects that consume energy and hinder performance.

These technologies, which contribute to higher performance and reduced power requirements, are the basis for enhancements to IBMs current processors and for IBMs future POWER4 Gigahertz processor. It is likely that these technologies will benefit many areas of system development.

1.3.3 IBM POWER4 Microprocessor

POWER4 is a new processor initiative from IBM. It is comprised of two 64-bit 1 GHz five issue superscalar cores that have a triple level cache hierarchy. It has a 10 GBps main memory interface with a 45 GBps multiprocessor interface. IBM is utilizing the 0.18 micron copper silicon-on-insulator technology in its manufacture. The targeted market is the Enterprise Server or servers in e-business. It is currently in the design stage.

1.3.4 Project Monterey

Project Monterey is the major UNIX initiative led by IBM, along with SCO and Intel. The Project Monterey UNIX operating system for IA-64 processors, code-named Monterey/64, is on track to be available when the first Intel Itanium processor-based hardware ships this year.

This past year, the Project Monterey initiative achieved a number of major goals, including becoming the first commercial UNIX operating system to run on the Itanium microprocessor. Project Monterey also recruited many leading hardware and software providers to join the initiative. Monterey/64 now has support from more original equipment manufacturers (OEMs) than any other operating system provider. Monterey/64 continues to gain ISV support and customer acceptance as the leading commercial UNIX operating system.

Launched October 26, 1998, the goal of Project Monterey is to establish a volume, commercial enterprise-class UNIX product line that runs across Intel IA-32, IA-64, and IBM POWER processors in systems that range from departmental to large data center servers. As part of this initiative, a UNIX operating system currently is being developed for the Intel IA-64 processor-based systems using IBM AIX, IBM NUMA-Q brand, and SCO UnixWare technologies. IBM is also providing the C and C++ compilers for building Monterey/64 based on IBM's compiler technology for AIX. The compilers are optimized for both the IA-64 and POWER architectures.

Chapter 2. RS/6000 Facts and Features Summary

The following section, taken from the brochure *Web Servers: RS/6000 and NUMA-Q Facts and Features*, G320-9878.

Figure 9 shows the RS/6000 Models that are the targets of this publication.



Figure 9. RS/6000 Systems Handbook Products in Focus

The following tables outline the important characteristics of the featured RS/6000 models.

Table 1. Facts and Features for Models 150, 170, and 270

RS/6000 Model	150	170	270
Machine Type	7043	7044	7044
Microprocessor Type	32-bit PowerPC 604e	64-bit POWER3-II	64-bit POWER3-II
# processors/system	1	1	1, 2, 3, or 4
Clock rates available (standard/option)	250/375 MHz	333/400 MHz	375 MHz
System Memory (Standard/Maximum)	128 MB/1 GB	256 MB/2 GB	256 MB/8 GB ^a
Data/instruction (L1) cache	32 KB/32 KB	64 KB/32 KB	64 KB/32 KB ^b
Level 2 (L2) cache	1 MB	1/4 MB	4 MB ^b
Reliability, Availability, Serviceability			
Service Processor	-	X	X
Hot-swappable disks (internal and external)	-	-	-
Dynamic CPU Deallocation ^j	-	-	X
Hot-plug slots	-	-	-
Redundant hot-plug power	-	-	-
Redundant hot-plug cooling	-	-	-
Capacity			
Slots available	5 PCI (32-bit)	6 PCI (4 32-bit + 2 64-bit)	5 PCI (2 32-bit+ 3 64-bit)
PCI bus speed	33 MHz	33/50 MHz	33/50 MHz
Disk/media bays	2/3	3/3	2/3
Minimum/maximum internal disk	9.1 GB/54.6 GB	9.1 GB/72.8 GB	9.1 GB/54.6 GB
Storage Interfaces			
SCSI-2 Fast/Wide	-	-	-
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra2 SCSI	X	X	X
PCI 3-Channel Ultra2 SCSI RAID	X	X	X
SSA Multi-Initiator RAID Enhanced Loop	-	-	-
SSA Advanced SerialRAID Plus	X	X	X
Gigabit Fibre Channel	-	-	-
Communications and Connectivity			
EIA RS232D/EIA RS422A	X	X	X
Token ring 4/16 Mbps	X	X	X
4-Port 10/100 Mbps Ethernet	X	X	X
Ethernet 10/100 Mbps	X	X	X
Gigabit Ethernet	-	X	X
FDDI 100 Mbps	X	X	X
ATM 155 Mbps	X	X	X
ISDN	X	X	X
SP System Attachment	-	-	-
ESCON Control Unit (host attach) ^k	-	-	-
ESCON Emulation (tape attach) ^k	-	-	-
HIPPI ^k	-	-	-
Digital Trunk Quad/Resource Adapter ^k	X/X	X/X	X/X
X.25 ^k	X	X	X
SDLC	X	X	X
BSC	X	X	X
Graphics Accelerators Available	GXT130P, 300P, 2000P, 3000P	GXT130P, 300P, 2000P, 3000P	GXT130P, 300P, 2000P, 3000P

Table 2. Facts and Features for Models B50, F50, and F80

RS/6000 Model	B50	F50	F80
Machine Type	7046	7025	7025
Microprocessor Type	32-bit PowerPC 604e	32-bit PowerPC 604e	64-bit RS64 III
# processors/system	1	1, 2, 3, or 4	1, 2, 4, or 6
Clock rates available (standard/option)	375 MHz	332 MHz	450 MHz (1, 2, 4), 500 MHz (6)
System Memory (Standard/Maximum)	128 MB/1 GB	128 MB/3 GB ^a	256 MB/16 GB ^a
Data/instruction (L1) cache	32 KB/32 KB	32 KB/32 KB ^b	128 KB/128 KB ^b
Level 2 (L2) cache	1 MB	256 KB ^b	2 MB(1), 4 MB(2, 4, 6) ^b
Reliability, Availability, Serviceability			
Service Processor	-	X	X
Hot-swappable disks (internal and external)	-	X	X
Dynamic CPU Deallocation ^j	-	-	X
Hot-plug slots	-	-	X
Redundant hot-plug power	-	-	O
Redundant hot-plug cooling	-	-	O
Capacity			
Slots available	2 PCI (32-bit)	7 PCI (5 32-bit +2 64-bit) + 2 PCI/ISA (32-bit)	10 PCI (64-bit)
PCI bus speed	33 MHz	33/50 MHz	33/66 MHz
Disk/media bays	2 ^h /2	18/4	14/3
Minimum/maximum internal disk	9.1 GB/54.6 GB	9.1 GB/345.8 GB	9.1 GB/254.8 GB
Storage Interfaces			
SCSI-2 Fast/Wide	-	X	X
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra2 SCSI	X	X	X
PCI 3-Channel Ultra2 SCSI RAID	X	X	X
SSA Multi-Initiator RAID Enhanced Loop	-	-	-
SSA Advanced SerialRAID Plus	X	X	X
Gigabit Fibre Channel	-	X	X
Communications and Connectivity			
EIA RS232D/EIA RS422A	X	X	X
Token ring 4/16 Mbps	X	X	X
4-Port 10/100 Mbps Ethernet	X	X	-
Ethernet 10/100 Mbps	X	X	X
Gigabit Ethernet	-	X	X
FDDI 100 Mbps	X	X	X
ATM 155 Mbps	X	X	X
ISDN	X	X	X
SP system attachment	-	-	-
ESCON Control Unit (host attach) ^k	-	X	X
ESCON Emulation (tape attach) ^k	-	X	X
HIPPI ^k	-	X	X
Digital Trunk Quad/Resource Adapter ^k	-/-	X/X	X/X
X.25 ^k	X	X	X
SDLC	X	X	X
BSC	X	X	X
Graphics Accelerators Available	GXT130P	GXT130P	GXT130P

Table 3. Facts and Features for Models H70 and H80

RS/6000 Model	H70	H80
Machine Type	7026	7026
Microprocessor Type	64-bit RS64 II	64-bit RS64 III
# processors/system	1, 2, 3, or 4	1, 2, 4, or 6
Clock rates available (standard/option)	340 MHz	450 MHz (1, 2, 4), 500 MHz (6)
System Memory (Standard/Maximum)	128 MB/8 GB ^a	256 MB/16 GB ^a
Data/instruction (L1) cache	64 KB/64 KB ^b	128 KB/128 KB ^b
Level 2 (L2) cache	4 MB ^b	2 MB (1), 4 MB (2, 4, 6) ^b
Reliability, Availability, Serviceability		
Service Processor	X	X
Hot-swappable disks (internal and external)	X	External only
Dynamic CPU Deallocation ^l	-	X
Hot-plug slots	-	X
Redundant hot-plug power	O	O
Redundant hot-plug cooling	X	X
Capacity		
Slots available	8 PCI (4 32-bit + 4 64-bit)	28 PCI (8 32-bit + 20 64-bit)
PCI bus speed	33/50 MHz	33/66 MHz
Disk/media bays	13/3	2/4
Minimum/maximum internal disk	9.1 GB/254.8 GB	0 GB/36.4 GB
Storage Interfaces		
SCSI-2 Fast/Wide	X	X
Ultra SCSI SE and Ultra SCSI Differential	X	X
PCI 2-Channel Ultra2 SCSI	X	X
PCI 3-Channel Ultra2 SCSI RAID	X	X
SSA Multi-Initiator RAID Enhanced Loop	-	-
SSA Advanced SerialRAID Plus	X	X
Gigabit Fibre Channel	X	X
Communications and Connectivity		
EIA RS232D/EIA RS422A	X	X
Token ring 4/16 Mbps	X	X
4-Port 10/100 Mbps Ethernet	X	-
Ethernet 10/100 Mbps	X	X
Gigabit Ethernet	X	X
FDDI 100 Mbps	X	X
ATM 155 Mbps	X	X
ISDN	X	X
SP system attachment	-	-
ESCON Control Unit (host attach) ^k	X	X
ESCON Emulation (tape attach) ^k	X	X
HIPPI ^k	X	X
Digital Trunk Quad/Resource Adapter ^k	X/X	X/X
X.25 ^k	X	X
SDLC	X	X
BSC	X	X
Graphics Accelerators Available	GXT130P	GXT130P

Table 4. Facts and Features for Models M80 and S80

RS/6000 Model	M80	S80
Machine Type	7026	7017
Microprocessor Type	64-bit RS64 III	64-bit RS64 III
# processors/system	2, 4, 6, or 8	6, 12, 18, or 24
Clock rates available (standard/option)	500 MHz	450 MHz
System Memory (Standard/Maximum)	1 GB/32 GB ^a	2 GB/64 GB ^a
Data/instruction (L1) cache	128 KB/128 KB ^b	128 KB/128 KB ^b
Level 2 (L2) cache	4 MB ^b	8 MB ^b
Reliability, Availability, Serviceability		
Service Processor	X	X
Hot-swappable disks (internal and external)	External only	X
Dynamic CPU Deallocation ^l	X	X
Hot-plug slots	X	-
Redundant hot-plug power	X	X
Redundant hot-plug cooling	X	X
Capacity		
Slots available	56 PCI (16 32-bit + 40 64-bit)	53 PCI (33 32-bit + 20 64-bit)
PCI bus speed	33/66 MHz	33 MHz
Disk/media bays	2/8	48/8
Minimum/maximum internal disk	0 GB/36.4 GB	9.1 GB/873.6 GB
Storage Interfaces		
SCSI-2 Fast/Wide	X	-
Ultra SCSI SE and Ultra SCSI Differential	X	X
PCI 2-Channel Ultra2 SCSI	X	X
PCI 3-Channel Ultra2 SCSI RAID	X	-
SSA Multi-Initiator RAID Enhanced Loop	-	-
SSA Advanced SerialRAID Plus	X	X
Gigabit Fibre Channel	X	X
Communications and Connectivity		
EIA RS232D/EIA RS422A	X	X
Token ring 4/16 Mbps	X	X
4-Port 10/100 Mbps Ethernet	-	-
Ethernet 10/100 Mbps	X	X
Gigabit Ethernet	X	X
FDDI 100 Mbps	X	X
ATM 155 Mbps	X	X
ISDN	X	X
SP system attachment	-	X
ESCON Control Unit (host attach) ^k	X	X
ESCON Emulation (tape attach) ^k	X	X
HIPPI ^k	X	-
Digital Trunk Quad/Resource Adapter ^k	X/-	-/-
X.25 ^k	X	X
SDLC	X	X
BSC	X	X
Graphics Accelerators Available	GXT130P	GXT130P

Table 5. Facts and Features for SP POWER3 nodes

RS/6000 Model	SPTM System (9076)^e		
Node Type	375 MHz POWER3 SMP Thin	375 MHz POWER3 SMP Wide	375 MHz POWER3 SMP High
Microprocessor Type	64-bit POWER3-II	64-bit POWER3-II	64-bit POWER3-II
# Min/Max of each node type per system	1/128 ^f	1/128 ^f	1/128 ^f
# processors/node	2 or 4	2 or 4	4, 8, 12, or 16
Clock rates available (standard/option)	375 MHz	375 MHz	375 MHz
System Memory (Standard/Maximum)	256 MB/8 GB ^a	256 MB/8 GB ^a	1 GB/64 GB ^a
Data/instruction (L1) cache	64 KB/32 KB ^b	64 KB/32 KB ^b	64 KB/32 KB ^b
Level 2 (L2) cache	8 MB ^b	8 MB ^b	8 MB ^b
Reliability, Availability, Serviceability			
Service Processor	X ^d	X ^d	X ^d
Hot-swappable disks (internal and external)	External only	External only	X ^g
Dynamic CPU Deallocation ^l	X	X	X
Hot-plug slots	-	-	-
Redundant hot-plug power	-	-	-
Redundant hot-plug cooling	-	-	-
Capacity			
Slots available	2 PCI (32-bit)	10 PCI (2 32-bit + 8 64-bit)	53 PCI (1 32-bit + 52 64-bit) ^g
PCI bus speed	33 MHz	33 MHz	33 MHz
Disk/media bays	2	4	2/26 ^g
Minimum/maximum internal disk	0 GB/36.4 GB	0 GB/109.2 GB	0 GB/946.4 GB ^g
Storage Interfaces			
SCSI-2 Fast/Wide	-	-	-
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra2 SCSI	X	X	X
PCI 3-Channel Ultra2 SCSI RAID	-	-	-
SSA Multi-Initiator RAID Enhanced Loop	-	-	-
SSA Advanced SerialRAID Plus	X	X	X
Gigabit Fibre Channel	X	X	X
Communications and Connectivity			
EIA RS232D/EIA RS422A	X	X	X
Token ring 4/16 Mbps	X	X	X
4-Port 10/100 Mbps Ethernet	X	X	X
Ethernet 10/100 Mbps	X	X	X
Gigabit Ethernet	X	X	X
FDDI 100 Mbps	X	X	X
ATM 155 Mbps	X	X	X
ISDN	-	-	-
SP system attachment	-	-	-
ESCON Control Unit (host attach) ^k	X	X	X
ESCON Emulation (tape attach) ^k	X	X	X
HIPPI ^k	X	X	X
Digital Trunk Quad/Resource Adapter ^k	X/X	X/X	X/X
X.25 ^k	X	X	X
SDLC	X	X	X
BSC	X	X	X
Graphics Accelerators Available	-	-	-

Notes:

X = Supported

O = Optionally Available

N/A = Not Applicable

a Shared memory

b Per processor

c Provides RAID 5 capability

d Using control workstation and PSSP software

e Node types can be intermixed on system

f Up to 512 available using special order

g With SP Expansion I/O units

h Third disk bay available via RPQ

j Requires firmware and AIX 4.3.3

k Requires additional software

Chapter 3. Workstations and Workgroup Servers

This chapter covers the following models:

- 7043-140 - Entry-Level Workstation or Entry Workgroup Server
- 7043-150 - Price/Performance Workstation or Entry Workgroup Server
- 7043-260 - 64-Bit Multiprocessor Workstation or 64-bit Workgroup Server
- 7044-170 - 64-Bit Design Workstation or 64-bit Workgroup Server
- 7044-270 - 64-Bit Multiprocessor Workstation or 64-bit Workgroup Server
- 7046-B50 - ISP/ASP Server Solution

The following models are discussed briefly for the purpose of comparison to more current models. They were withdrawn from IBM marketing on May 8, 2000, March 19, 1999, and January 8, 1998 respectively.

- 7025-F40 - Workstation or Workgroup Server
- 7043-240 - Workstation or Workgroup Server
- 7024-E30 - Workgroup Server

Workstations are desktide or desktop machines that excel in numerically and graphically intensive applications that are usually run for a single user.

Workgroup Servers are desktide machines that contain a substantial amount of storage and other features used to support the clients for a midsize company or a medium- to large-size department, depending on the chosen application. All RS/6000 workstations can be equipped with features that allow them to be used as a workgroup server.

3.1 Model History and Uses

The RS/6000 43P and 44P Series is a line of full-function desktop systems that deliver performance and expansion capabilities at affordable prices. These systems are versatile and perform well as either workstations or entry workgroup servers.

The following are brief descriptions of the uses and histories of the machine type 7043 RS/6000 Models 140, 150, and 260, machine type 7044 for the 170 and 270, and machine type 7026 Model B50.

Note

The 7043 Models 140, 150, and 260 are referred to as 43P machines because they are the follow-on products to the 7248 Model 43P. The 7044 Models 170 and 270 are referred to as 44Ps.

3.1.1 43P Model 140

The 7043 family was launched on November 8, 1996 with the availability of the Model 140. The Model 140, a uniprocessor system, is an ideal workgroup server for running small business and departmental applications. It is also an excellent solution for demanding 2D, but is suitable for midrange 3D graphics applications. The Model 140 will be withdrawn from marketing December 29, 2000.

3.1.2 43P Model 150

The Model 150 became available on October 23, 1998 as the high-end compliment to the Model 140. The Model 150 is a uniprocessor system that provides enhanced performance over its predecessor, the Model 140, by utilizing a faster processor, an enhanced memory controller, advanced graphics, and integrated 10/100 Ethernet and Ultra SCSI controllers on the planar. It performs well as a high-function 3D graphics workstation, a departmental or small business workgroup server, or an e-business server.

3.1.3 44P Model 170

On February 7, 2000, IBM announced its first 64-bit uniprocessor system based on the copper chip POWER3-II technology, the Model 170. A new hardware enclosure and an additional adapter slot add value in addition to the performance. The Model 170 is a compact and reliable system that offers the high performance and expandability you need to help grow your business. It is an excellent choice for a workgroup, technical, or e-business server and is ideally suited for 64-bit SMP application development and testing.

3.1.4 43P Model 260

On October 05, 1998, the Model 260 was announced as the successor to the SMP-enabled Model 240. The Model 260 provides significant performance enhancements over the Model 240 with its 64-bit capability, enhanced processor, enhanced memory controller, advanced graphics, and planar-integrated Ethernet and Ultra SCSI controllers. With its enhanced floating point capability, the Model 260 is designed for memory and compute-intensive analysis and can be used as a single-seat mechanical

computer-aided design (MCAD) and analysis solution. It also provides excellent performance, expandability, and reliability for departments and small businesses as an e-business or entry workgroup server or as a cost-effective development platform for developing and testing applications that will run on larger RS/6000 systems. The dual processing power of the Model 260 and its small package make it an excellent solution for Internet service providers and customers that need a stand-alone Internet server.

3.1.5 44P Model 270

The Model 270 was announced on February 7, 2000 and extends IBMs line of powerful and affordable workstations using 64-bit copper chip POWER3-II technology. Designed for technical and graphics users who require high-performance and symmetric multiprocessing (SMP) capability, the Model 270 workstation lowers the cost of high-end design and analysis solutions while significantly raising levels of performance. The Model 270 is also a very powerful Web and e-commerce server.

3.1.6 Model B50

The model B50 was announced on September 13, 1999. The Model B50 is one of IBMs first server offerings to specifically target the Internet and application server provider segments where small footprint and high-performance Web and application serving are critical business requirements and is focused on ISP and ASP environments where database transaction, e-commerce, and dynamic content demand this level of price and performance. It is one of the first Linux-enabled platforms in the RS/6000 product line. With a slide out design, balanced mechanical chassis, ergonomically-located grab handles, and easy access to the adapters and disk drives without the use of tools, expansion and replacement are simple tasks. Aggressive pricing and an integrated AIX software license make the Model B50 the most affordable new model in the RS/6000 product line.

3.1.7 Upgrade Paths

The following Workstations and Workgroup Server upgrade paths are currently available:

- 7025-F40 -> 7025-F50
- 7043-260 -> 7043-270

No upgrade paths are currently available for the 7043-140, 7043-150, and 7044-170 models. Depending on your current configuration, you may be able to upgrade your processor on a 43P Model 140 up to a clockrate of 233 MHz or 332 MHz or add a second processor to a 43P Model 260. Similarly, you can

increase the processor speed of the 44P Model 170 from 332 MHz to 400 MHz and the 44P Model 270 from the default 1-way to a 3-way or from a 2-way to a 3-way or 4-way.

3.2 RS/6000 43P 7043 Model 140 Overview

Check with your IBM sales representative for the latest availability of the Model 140.

The following sections list the standard and optional features of the Model 140 as shown in Figure 10 on page 28.



Figure 10. Model 7043-140 - Front View

3.2.1 43P Model 140 Standard Configuration

Table 6 provides the standard configuration for the Model 140.

Table 6. Model 140 Standard Configuration

Model 140 Standard Configuration and Standard Features	
Microprocessor	233 MHz PowerPC 604e

Model 140 Standard Configuration and Standard Features	
Level 1 (L1) cache	32 KB data/32 KB instruction
Level 2 (L2) cache	1 MB
RAM (minimum)	64 MB ECC DIMM
Memory bus width	64-bit
Integrated ports	Tablet, keyboard, mouse, Ethernet (thick + twisted pair), SCSI-2 F/W, serial (two), parallel, and stereo audio
Internal disk drive	9.1 GB SCSI-2
Disk/media bays	Five (one disk and one media available)
Expansion slots	Five (three 32-bit PCI + two shared PCI/ISA)
PCI bus width	32-bit
Memory slots	six
CD-ROM drive	20X (Max) SCSI-2 CD-ROM drive
Diskette drive	1.44 MB 3.5-inch diskette drive
SCSI adapters	Integrated SCSI-2 F/W adapter
Service processor	No
AIX operating system version	AIX 4.1.5, AIX 4.2.1, AIX 4.3.1, or later
System dimensions and weight	6.5" H x 16.5" W x 18.1" D (165 mm x 420 mm x 460 mm); 32 lbs (14.5 kg)
Warranty	24 x 7, on-site for one year (limited) at no additional cost.

3.2.2 43P Model 140 System Expansion

Table 7 shows the possible processor, storage, and memory expansion configurations.

Table 7. Model 140 System Expansion

Model 140 System Expansion	
Processor upgrade	332 MHz PowerPC 604e
RAM	Up to 768 MB

Model 140 System Expansion	
Internal disk storage	Up to 54.6 GB
External disk storage	For additional information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.2.3 43P Model 140 Optional Features

The following section outlines the optional features of the Model 140. *Optional features* are defined as optional internal devices and internal adapters that can be configured on the RS/6000 Model 140.

Table 8 lists the optional internal features of the Model 140 and indicates whether each feature is available or withdrawn.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.
- W** These features are withdrawn. At one point in time, they were available for this product.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 8. Model 140 Optional Features

Feature Code	Description	Status
Processors		
4318	233 MHz Pluggable Processor	A
4347	332 MHz Pluggable Processor	A
Memory		
9075	Base 64 MB DRAM DIMM 60 ns, 168 PIN	W
4113	32 MB DRAM DIMM 60 ns, 168 PIN	A
4114	64 MB DRAM DIMM 60 ns, 168 PIN	A
4115	128 MB DRAM DIMM 60 ns, 168 PIN	A
Internal Disk Drives		
2501	9.1 GB 1" Ultra SCSI 7200 RPM	W

Feature Code	Description	Status
2502	9.1 GB 1" Ultra SCSI 7200 RPM	A
2503	18.2 GB 1" Ultra SCSI 7200 RPM	A
2504	18.2 GB 1" Ultra SCSI 7200 RPM	W
2908	9.1 GB 1" Ultra SCSI	W
2909	18.2 GB 1" Ultra SCSI	W
3114	18.2 GB 1" Ultra SCSI	W
Internal Tape Drives		
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2619	20X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
2841	GXT300P	A
2823	GXT2000P	A
Multimedia Adapters		
2639	Ultimedia Video Capture (Short)	A
SCSI Adapters		
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
SSA Adapters		
6222	SSA Fast-Write Cache Option	A
6230	Advanced SerialRAID Plus Adapter	A
6231	128 MB Cache Option for 6225	A
6235	Advanced SerialRAID Fast Write Cache	A
Asynchronous Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A

Feature Code	Description	Status
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
2998	Turboways 25 PCI	A
Token-ring adapters		
2920	4/16 Mbps token-ring adapter	W
4959	4/16 Mbps token-ring adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	A
2985	Ethernet BNC / RJ-45	A
2987	Ethernet AUI / RJ-45	A
4951	4-Port 10/100 Base-TX Ethernet	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
X.25 Adapter		
2962	2-Port Multiprotocol PCI	A
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A

Feature Code	Description	Status
3753	Options Library (hard copy)	W
6041	Mouse, 3 button, white	A
8741	Mouse, 3 button, black	A
8227	Security strap	A

3.2.4 43P Model 140 Configuration Notes

It is important to keep the following in mind when configuring the Model 140:

- When configuring a third disk drive on the Model 140, you must mount the drive in a media bay, which requires the media bay disk drive mounting kit (# 6509).
- PCI and graphics adapters are subject to placement rules, which are tabulated in Appendix D.5, "7043 Model 140 Adapter Placement Guide" on page 664.
- A total of six memory DIMMs can be configured in capacities of 32 MB, 64 MB, and 128 MB. Memory DIMMs do not have to be installed in pairs, and sizes can be mixed.
- For a complete list of system configuration limitations, refer to the Limitations section of the Sales Manual.

3.2.5 43P Model 140 Publications

Table 9 provides the publications available for the Model 140.

Table 9. Publications for the Model 140

Order Number	Title
SA38-0510	7043 43P Series Setup Instructions
SA38-0511	7043 43P Series Users Guide
SA23-2690 (USA customers only)	Customer Support Information
SA23-2652	System Unit Safety Information
SA38-0538	PCI Adapter Placement Reference Guide
Customer Installable Options Library CD-ROM ¹	
Warranty Booklet (USA customers only)	
Notes: 1 The CD-ROM is not orderable. It is shipped only at manufacturing. No form number is available.	

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

3.3 RS/6000 43P 7043 Model 150 Overview

The following sections discuss the minimum configuration and optional features of the RS/6000 Model 150, which is shown in Figure 11.




Figure 11. Model 7043-150 with Peripherals

3.3.1 43P Model 150 Minimum Configuration

Table 10 lists the characteristics of a Model 150 minimum configuration.

Table 10. Model 150 Minimum Configuration

Model 150 Minimum Configuration	
Microprocessor	250 MHz or 375 MHz PowerPC 604e
Level 1 (L1) cache	32 KB data/32 KB instruction
Level 2 (L2) cache	1 MB
RAM (minimum)	128 MB SDRAM
Memory bus width	64-bit

Model 150 Minimum Configuration	
Integrated ports	Tablet, keyboard, mouse, 10/100 Ethernet, Ultra SCSI, serial (two), parallel, and stereo audio
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Five (one disk and one media bay available)
Expansion slots	Five PCI
PCI bus width	32-bit
Memory slots	Four
CD-ROM drive	32X (Max) SCSI-2 CD-ROM
Diskette drive	1.44 MB 3.5-inch diskette drive
Service Processor	No
AIX operating system version	AIX 4.3.2, or later
System dimensions and weight	6.5" H x 16.5" W x 18.1" D (165 mm x 420 mm x 460 mm); 32 lbs (14.5 kg)
Warranty	24 x 7, on-site for one year (limited), at no additional cost

3.3.2 43P Model 150 System Expansion

Table 11 shows the possible storage and memory expansion configurations.

Table 11. Model 150 System Expansion

Model 150 System Expansion	
RAM	Up to 1 GB of SDRAM
Internal disk storage	Up to 54.6 GB
External disk storage	For additional information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.3.3 43P Model 150 Optional Features

The following section outlines the optional features of the Model 150. *Optional features* are defined as optional internal devices and internal adapters that can be configured on the RS/6000 Model 150.

Table 12 lists the optional internal features of the Model 150.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, monitors, are not included.

Table 12. Model 150 Optional Features

Feature Code	Description	Status
Processor		
4348	375 MHz 604e Processor	A
4363	250 MHz 604e Processor	A
Memory		
4150	128 MB SDRAM DIMM	A
4169	256 MB SDRAM DIMM	A
Internal Disk Drives		
2908	9.1 GB 1" Ultra SCSI	A
2909	18.2 GB 1" Ultra SCSI	A
Internal Tape Drives		
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
2841	GXT300P	A
2823	GXT2000P	A
2825	GXT3000P	A
Multimedia Adapters		
2639	Ultimedia Video Capture/S (Short)	A

Feature Code	Description	Status
SCSI Adapters		
2494	Ultra SCSI 3-Channel RAID	A
6205	Dual Channel Ultra2 SCSI	A
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
SSA Adapters		
6222	SSA Fast-Write Cache	A
6230	Advanced SerialRAID Plus Adapter	A
6231	128 MB DRAM Option Card for 6225	A
6235	32 MB Fast-Write Cache Option Card	A
Asynchronous Adapters		
2943	8-Port Asynchronous Adapter EIA-232/RS-422	A
2944	128-Port Asynchronous Controller	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital Trunk	A
6311	ARTIC960RxF Digital Trunk Resource	A
ATM Adapters		
2963	Turboways 155 PCI UTP	A
2988	Turboways 155 PCI MMF ATM	A
2998	TURBOWAYS 25 ATM PCI	A
Token-ring adapters		
4959	Token-ring adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	A
4951	4-Port 10/100 Base-TX Ethernet PCI	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A

Feature Code	Description	Status
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A
6041	Mouse, 3 button, white	A
8227	Security strap	A
8741	Mouse, 3 button, black	A

3.3.4 43P Model 150 Configuration Notes

It is important to keep the following notes in mind when configuring the Model 150:

- When configuring a third disk drive on the Model 150, you must mount the drive in a media bay, which requires the media bay disk drive mounting kit (# 6509).
- PCI and graphics adapters are subject to placement rules, which are tabulated in Appendix D.6, “7043 Model 150 Adapter Placement Guide” on page 667.
- Up to four memory DIMMs (128 MB, 256 MB) can be installed. These DIMMs do not have to be installed in pairs.
- For a complete list of system configuration limitations, refer to the Limitations section of the Model 150 within the Sales Manual.

3.3.5 43P Model 150 Publications

Table 13 lists the publications that are for the Model 150.

Table 13. Publications for the Model 150

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

3.4 RS/6000 44P 7044 Model 170 Overview

The following sections outline the standard and optional features of the Model 170 as shown in Figure 12.



Figure 12. Model 7044-170 with Peripherals

3.4.1 44P Model 170 Minimum Configuration

Table 14 lists the characteristics of a Model 170 minimum configuration.

Table 14. Model 170 Minimum Configuration

Model 170 Minimum Configuration	
Microprocessor	333 MHz POWER3-II
Level 1 (L1) Cache	64 KB data/32 KB instruction
Level 2 (L2) Cache	1 MB
RAM (minimum)	256 MB SDRAM
Memory bus width	128-bit
Integrated ports	Tablet, keyboard, mouse, 10/100 Mbps Ethernet (thick and twisted), internal Ultra SCSI, external Ultra2 SCSI, serial (two), parallel and stereo audio
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Six (three available)
Expansion slots	Six PCI
PCI bus width	32- and 64-bit
Memory slots	Four
CD-ROM drive	32X (Max) SCSI-2 CD-ROM drive
Diskette drive	1.44 MB 3.5-inch diskette drive
Service Processor	Standard (on planar)
AIX operating system version	AIX 4.3.3 (unlimited user license)
System dimension and weight	19.25" H x 7.9" W x 20.25" D (490 mm x 200 mm x 515 mm); 39 lbs. (17.7 kg)
Warranty	On-site for one year (limited) at no additional cost

3.4.2 44P Model 170 System Expansion

Table 15 shows the possible processor, storage, and memory expansion configurations.

Table 15. Model 170 System Expansion

Model 170 System Expansion	
Microprocessor	400 MHz POWER3-II

Model 170 System Expansion	
Level 2 (L2) Cache	4 MB
RAM (memory)	Up to 2 GB
Internal disk storage	Up to 145.6 GB
External disk storage	Options include: IBM 2104 Expandable Storage Plus (Ultra2 SCSI), IBM 7133 Serial Disk System (SSA). For additional information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.4.3 44P Model 170 Optional Features

The following section outlines the optional features of the Model 170. *Optional features* are defined as optional internal devices and internal adapters that can be configured on the RS/6000 Model 170.

Table 16 lists the optional internal features of the Model 170.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 16. Model 170 Optional Features

Feature Code	Description	Status
Processor		
4349	333 MHz POWER3-II Processor Card with 1 MB L2 Cache	A
4360	400 MHz POWER3-II Processor Card with 4 MB L2 Cache	A
Memory		
4110	256 MB (2 X 128 MB) Dimms, 200 pin 10ns SDRAM	A
4120	512 MB (2 X 256 MB) Dimms, 200 pin 10ns SDRAM	A
4121	1024 MB (2 X 512 MB), Dimms, 200 pin 10ns SDRAM	A

Feature Code	Description	Status
Internal Disk Drives		
2908	9.1 GB 1" Ultra SCSI	A
3027	9.1 GB 10K RPM Ultra SCSI	A
2909	18.2 GB 1" Ultra SCSI	A
3102	18.2 GB 10K RPM Ultra SCSI	A
3119	36.4 GB 10K RPM Ultra SCSI	A
Internal Tape Drive		
6154	20/40 GB 16-bit 8 mm Internal Tape Drive (White Bezel)	A
6156	20/40 GB 16-bit 8 mm Internal Tape Drive (Black Bezel)	A
6159	12/24 GB 4 mm Internal Tape Drive	A
Internal CD-ROM Drive		
2624	SCSI-2 CD-ROM Drive	A
Graphics Accelerators		
2830	POWER GXT130P	A
2841	POWER GXT300P	A
2823	POWER GXT2000P	A
2825	POWER GXT3000P	A
Multimedia Adapter		
2639	Ultimedia Video Capture Adapter/S (Short)	A
SCSI Adapters		
2494	PCI 3-Channel Ultra2 SCSI RAID Adapter	A
6204	Universal Differential Ultra SCSI Adapter	A
6205	Dual Channel Ultra2 SCSI Adapter	A
6206	Single-Ended Ultra SCSI Adapter	A
SSA Adapters		
6230	Advanced SerialRAID Plus Adapter	A

Feature Code	Description	Status
6231	128 MB DRAM Option Card	A
6235	32 MB Fast-Write Cache Option Card	A
Async Adapters		
2943	8-Port Asynchronous Adapter EIA-232/RS-422	A
2944	128-Port Asynchronous Controller	A
ARTIC Adapters		
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	A
Digital Trunk Adapters		
6310	IBM ARTIC960RxD Quad Digital Trunk PCI Adapter	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM Adapter	A
2988	Turboways 155 PCI MMF ATM Adapter	A
2998	Turboways 25 ATM PCI Adapter	A
Token-ring adapters		
4959	4/16 Mbps Token-ring adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	A
2969	Gigabit Ethernet - SX PCI Adapter	A
4951	4-Port 10/100 Mbps Ethernet PCI Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI Adapter	A
FDDI		
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	A
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	A
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	A
ISDN Adapters		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter	A

Feature Code	Description	Status
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A
6041	Mouse, 3 button, white	A
8227	Security Tie Down Strap	A
8741	Mouse, 3 button, black	A

For more information on the adapters, refer to Chapter 9, “Communications and Storage I/O Adapters” on page 323, or see the *PCI Adapter Placement Reference*, SA38-0538.

3.4.4 44P Model 170 Configuration Notes

It is important to keep the following notes in mind when configuring the Model 170:

- Bay 2 can accommodate a second CD-ROM, tape drive, or other media device.
- A 9.1 GB (# 2908 or # 3027), 18.2 GB (# 2909), or 36.4 GB (# 3119) hard disk can be installed in Bay 2 with a mounting kit (# 6560).
- A minimum of 256 MB of memory is required, the maximum being 2 GB of system memory. One to four memory DIMMs (256 MB or 512 MB) can be installed (the DIMMs must be installed in pairs).
- PCI and graphics adapters are subject to placement rules, which are tabulated in Appendix D.7, “7044 Model 170 Placement Guide” on page 671.
- For a complete list of system configuration limitations, refer to the Limitations section of the Model 170 within the Sales Manual.
- The 10 K RPM 18.2 disk drive is not available for Bay 2 (# 3102).
- Whenever using an async terminal as a console device on 7043 and 7044 systems; a DB-9 to DB-25 convertor is required. This necessary cable is not shipped with all models that require it.
 - The convertor cable automatically ships with:
 - 7043-140
 - 7043-150
 - 7043-260

- The convertor cable must be ordered separately for:
 - 7044-170
 - 7044-270
- (# 3925) Serial Port Converter Cable 9-pin to 25-pin, pigtail
- (# 3926) Async Printer/Terminal Cable 9-pin to 25-pin, 4 m length

3.4.5 44P Model 170 Publications

The publications, listed in Table 17, are available for the Model 170 (additional copies are available):

Table 17. Publications for the Model 170

Order Number	Title
SA38-0559	<i>44P Series Model 170 Users Guide</i>
SA38-0561	<i>44P Series Model 170 Setup Instructions</i>
SA38-2652	<i>System Unit Safety Information</i>
SA38-2690	<i>Customer Support Information (US customers only)</i>
SA38-0538	<i>PCI Adapter Placement Reference Guide</i>

Table 18 provides the additional publications that are available. To order, contact your IBM representative.

Table 18. Additional Publications for the Model 170

Order Number	Title
SA38-0560	<i>44P Series Model 170 Service Guide</i>
SA38-0509	<i>Diagnostics Information for Multiple Bus Systems</i>
SA38-0516	<i>RS/6000 Adapters, Devices and Cables Information for Multiple Bus Systems</i>

3.5 RS/6000 43P 7043 Model 260 Overview

The following sections list the minimum configuration and optional features of the Model 260 shown in Figure 13.



Figure 13. Model 7043-260 with Peripherals and Spaceball

3.5.1 43P Model 260 Minimum Configuration

Table 19 provides the minimum configuration of a Model 260.

Table 19. Model 260 Minimum Configuration

Model 260 Minimum Configuration	
Microprocessor	200 MHz POWER3
Level 1 (L1) cache	64 KB data/32 KB instruction per processor
Level 2 (L2) cache	4 MB per processor
RAM (minimum)	256 MB SDRAM
Memory bus width	128-bit

Model 260 Minimum Configuration	
Integrated ports	Tablet, keyboard, mouse, Ethernet 10/100, Ultra SCSI, serial (two) parallel, and stereo audio
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Five (one disk and one media bay available)
Expansion slots	Five (three 32-bit PCI + two 64-bit PCI)
PCI bus width	64-bit
Memory slots	Two memory card slots, 16 DIMM slots per card
CD-ROM drive	32X (Max) SCSI-2 CD-ROM
Diskette drive	1.44 MB 3.5-inch diskette drive
Service Processor	Standard
AIX operating system version	AIX 4.2.1, AIX 4.3.2 or later
System dimensions and weight	24.0" H x 13.3" W x 28.1" D (610 mm x 340 mm x 713 mm); 80/97 lbs (36.2/43.9 kg)
Warranty	24 x 7, on-site for one year (limited), at no additional cost

3.5.2 43P Model 260 System Expansion

Table 20 shows the possible processor, storage, and memory expansion configurations.

Table 20. Model 260 System Expansion

Model 260 System Expansion	
Processor	1- to 2-way SMP
RAM	Up to 8 GB of SDRAM
Internal disk storage	Up to 109.2 GB
External disk storage	For information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.5.3 43P Model 260 Supported Optional Features

The following section outlines the optional features of the Model 260. *Optional features* are defined as optional internal devices, external adapters, and external subsystems that can be configured on the RS/6000 Model 260.

Table 21 lists the optional internal features of the Model 260.

The status of a feature is indicative of these qualifications:

A Indicates features that are available and orderable on the specified models.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors are not included.

Table 21. Model 260 Optional Features

Feature Code	Description	Status
Processor		
4342	200 MHz POWER3 Processor	A
Memory		
4107	64 MB (2x32 MB) SDRAM DIMMs	A
4110	256 MB (2x128 MB) SDRAM DIMMs	A
4119	512 MB (2x256 MB) SDRAM DIMMs	A
4098	Memory Expansion Feature (16 Position)	A
Internal Disk Drives		
2908	9.1 GB 1" Ultra SCSI	A
3029	9.1 GB Ultra SCSI Enhanced	A
3027	9.1GB 10K RPM Ultra SCSI	A
2906	9.1 GB 10K RPM Ultra SCSI Enhanced	A
2909	18.2 GB 1" Ultra SCSI	A
3103	18.2 GB Ultra SCSI Enhanced	A
3102	18.2 GB 10K RPM Ultra SCSI	A
3110	18.2 GB 10K RPM Ultra SCSI Enhanced	A
3119	36.4 GB 10K RPM Ultra SCSI	A
3128	36.4 GB 10K RPM Ultra SCSI Enhanced	A
6505	Hot Swap Carrier for 16-bit SCSI	A
6508	Media Bay Disk Drive Mounting Kit	A

Feature Code	Description	Status
Internal Tape Drives		
6156	20/40 GB 8 mm Internal Tape Drive (Black Bezel)	A
6159	12/24 GB 4 mm Internal Tape Drive	A
Internal CD-ROMs		
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
2841	GXT300P	A
2823	GXT2000P	A
2825	GXT3000P	A
Multimedia Adapters		
2639	Ultimedia Video Capture Adapter/S (Short)	A
SCSI Adapters		
2494	Ultra SCSI 3-Channel RAID	A
6205	Dual Channel Ultra2 SCSI Adapter	A
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
SSA Adapters		
6222	SSA Fast-Write Cache	A
6230	Advanced SerialRAID Plus Adapter	A
6231	128 MB DRAM Option Card	A
6235	32 MB Fast-Write Cache Option Card	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Multiprotocol PCI	A

Feature Code	Description	Status
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital Trunk PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
2998	Turboways 25 ATM PCI	A
Token-ring adapters		
4959	4/16 Mbps token-ring adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	A
2969	Gigabit Ethernet - SX PCI Adapter	A
4951	4-Port 10/100 Mbps Ethernet PCI Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI Adapter	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	A
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	A
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter	A
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A
6041	Mouse, 3 button, white	A
8741	Mouse, 3 button, black	A

For more information on the adapters, refer to Chapter 9, “Communications and Storage I/O Adapters” on page 323, or see the *PCI Adapter Placement Reference*, SA38-0538.

3.5.4 43P Model 260 Configuration Notes

It is important to keep the following notes in mind when configuring the Model 260:

- When configuring a third disk drive on the Model 260, you must mount the drive in a media bay, which requires the media bay disk drive mounting kit (# 6508).
- PCI and graphics adapters are subject to placement rules, which are tabulated in Appendix D.8, "7043 Model 260 Adapter Placement Guide" on page 674.
- One or two memory cards may be installed. The first memory card must be installed in memory card slot 2. Each memory card can support eight pairs of DIMMs. Memory DIMMs must be ordered and installed in pairs on the memory cards. They must be installed starting at the bottom of each card (card slot J1 and J2).
- For a complete list of system configuration limitations, refer to the Limitations section of the Model 260 within the Sales Manual.

3.5.5 43P Model 260 Publications

Table 22 provides the common publications for the Model 260:

Table 22. Publications for the Model 260

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

Table 23 provides the additional publications that are available. To order, contact your IBM representative.

Table 23. Additional Publications for the Model 260

Order Number	Title
SA38-0554	<i>7043 Model 260 Service Guide</i>
SA38-0509	<i>RS/6000 Diagnostics Information for Multiple Bus Systems</i>

Order Number	Title
SA38-0516	<i>RS/6000 Adapters, Devices and Cables Information for Multiple Bus Systems</i>

3.6 RS/6000 44P 7044 Model 270 Overview

This section discusses the 7044-270 and the 7043-270 systems. The 7043-270 is the upgrade to a Model 270 from the 43P-260 system. Figure 14 shows the Model 270, and the following sections list the minimum configuration and optional features.




Figure 14. Model 7044-270 with Peripherals

3.6.1 44P Model 270 Minimum Configuration

Table 24 provides the minimum configuration of a Model 270.

Table 24. Model 270 Minimum Configuration

Model 270 Minimum Configuration	
Microprocessor	375 MHz POWER3-II

Model 270 Minimum Configuration	
Level 1 (L1) cache	64 KB data/32 KB instruction per processor
Level 2 (L2) cache	4 MB per processor
RAM (minimum)	256 MB SDRAM
Memory bus width	64-bit
Integrated ports	Tablet, keyboard, mouse, Ethernet 10/100, Ultra SCSI, serial (two), parallel, and stereo audio
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Five (one disk and one media bay available)
Expansion slots	Five (three 32-bit PCI + two 64-bit PCI)
PCI bus width	32-bit and 64-bit
Memory slots	Two memory card slots, 16 DIMM slots per card
CD-ROM drive	32X (Max) SCSI-2 CD-ROM
Diskette drive	1.44 MB 3.5-inch diskette drive
Service Processor	Standard
AIX operating system version	AIX 4.3.3 (unlimited user license)
System dimensions and minimum/maximum weight	24.0" H x 13.3" W x 28.1" D (610 mm x 340 mm x 713 mm); 81/99 lbs (37/45 kg)
Warranty	24 x 7, on-site for one year (limited), at no additional cost

3.6.2 44P Model 270 System Expansion

Table 22 shows the possible storage and memory configurations.

Table 25. Model 270 System Expansion

Model 270 System Expansion	
Processor	2, 3, or 4-way SMP 375 MHz POWER3-II
RAM	Up to 8 GB of SDRAM
Internal disk storage	Up to 109.2 GB

Model 270 System Expansion	
External disk storage options	Options include: 2104 Expandable Storage Plus, IBM 2105 Enterprise Storage Server, IBM 7133 SSA Disk Subsystem. For additional information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.6.3 44P Model 270 Supported Optional Features

The following section outlines the optional features of the Model 270. *Optional features* are defined as optional internal devices, external adapters, and external subsystems that can be configured on the RS/6000 Model 270.

Table 26 lists the optional internal features of the Model 270 and indicates whether each feature is available or supported.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.
- S** Indicates a supported device in this model. It cannot be ordered as a new feature but can be installed as part of an upgrade.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 26. Model 270 Optional Features

Feature Code	Description	Status
Processor		
4361	1-Way 375 MHz POWER3-II Processor Card	A
4362	2-Way 375 MHz POWER3-II Processor Card	A
Memory		
4110	256 MB (2x128 MB) SDRAM DIMMs	A
4120	512 MB (2X256MB) SDRAM DIMMs	A
4098	Memory Expansion Feature (16 POS)	A
Internal Disk Drives		
2900	4.5 GB Ultra SCSI	S
3028	4.5 GB Ultra SCSI Enhanced	S

Feature Code	Description	Status
2906	9.1 GB 10K RPM Ultra SCSI Enhanced	A
2908	9.1 GB 1" Ultra SCSI	A
3027	9.1 GB 10K RPM Ultra SCSI	A
3029	9.1 GB Ultra SCSI Enhanced	A
2909	18.2 GB 1" Ultra SCSI	A
3102	18.2 GB 10K RPM Ultra SCSI	A
3103	18.2 GB Ultra SCSI Enhanced	A
3110	18.2 GB 10K RPM Ultra SCSI Enhanced	A
3119	36.4 GB 10K RPM Ultra SCSI	A
3128	36.4 GB 10K RPM Ultra SCSI Enhanced	A
6505	Hot Swap Carrier for 16-bit SCSI	A
6508	Media Bay Disk Drive Mounting Kit	A
Internal Tape Drives		
6154	20/40 GB 8mm Internal Tape Drive (White Bezel)	A
6156	20/40 GB 8mm Internal Tape Drive (Black Bezel)	A
6159	12/24 GB 4mm Internal Tape Drive	A
Internal CD-ROMs		
2624	CD-ROM	A
Graphics Accelerators		
2838	GXT120P	S
2851	GXT250P	S
2852	GXT255P	S
2830	GXT130P	A
2841	GXT300P	A
2823	GXT2000P	A
2825	GXT3000P	A

Feature Code	Description	Status
Multimedia Adapters		
2639	Ultimedia Video Capture/S (Short)	A
SCSI Adapters		
2494	Ultra SCSI 3-Channel RAID	A
6204	PCI Universal Differential Ultra SCSI	A
6205	Dual Channel Ultra2 SCSI Adapter	A
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	S
SSA Adapters		
6225	Advanced SerialRAID Plus Adapter	S
6215	Advanced SerialRAID Plus Adapter	S
6230	Advanced SerialRAID Plus Adapter	A
6231	128 MB DRAM Option Card	A
6222	SSA Fast-Write Cache Option	A
6235	32 MB Fast-Write Cache Option Card	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2948	ARTIC960Hx 4-Port T1 PCI	S
2947	ARTIC960Hx 4-Port Multiprotocol PCI	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital Trunk PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
2998	Turboways 25 ATM PCI	A

Feature Code	Description	Status
Token-ring adapters		
2920	4/16 Mbps token-ring adapter	S
4959	4/16 Mbps token-ring adapter	A
Ethernet Adapters		
2968	10/100 Mbps Ethernet PCI	A
2969	Gigabit Ethernet - SX PCI	A
4951	4-Port 10/100 Mbps Ethernet PCI Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A
6041	Mouse, 3 button, white	A
8741	Mouse, 3 button, black	A

For more information on the adapters, refer to Chapter 9, “Communications and Storage I/O Adapters” on page 323, or see the *PCI Adapter Placement Reference*, SA38-0538.

3.6.4 44P Model 270 Dynamic Processor Deallocation

A high availability feature is dynamic processor deallocation. This feature allows the system to completely isolate a processor that has been determined to be operating below an acceptable level of reliability. To accomplish this, the system monitors processor status in real-time to determine when the number of recoverable errors has exceeded a specified threshold. When the threshold

is exceeded, the system sends a report to the operating environment that a processor is unreliable. The AIX operating environment then re-routes all interrupts and processes to other processors and disables the unreliable processor complex without interrupting end user services. Firmware level 991216 must be installed on the Model 270. The location of the URL that microcode or firmware can be downloaded from is:

<http://www.austin.ibm.com/support/micro/download.html>

3.6.5 44P Model 270 Configuration Notes

It is important to keep the following notes in mind when configuring the Model 270:

- When configuring a third disk drive on the Model 270, you must mount the drive in a media bay, which requires the media bay disk drive mounting kit (# 6508).
- PCI and graphics adapters are subject to placement rules, which are tabulated in Appendix D.9, “7043/7044 Model 270 Adapter Placement Guide” on page 677.
- One or two memory cards may be installed. The first memory card must be installed in memory card slot 2. Each memory card can support eight pairs of DIMMs. Memory DIMMs must be ordered and installed in pairs on the memory cards. They must be installed starting at the bottom of each card (card slot J1 and J2).
- For a complete list of system configuration limitations, refer to the Limitations section of the Model 270 within the Sales Manual.
- The 7044-270 does not currently support the gigabit fibre channel adapter.
- Whenever using an async terminal as a console device on 7043 and 7044 systems, a DB-9 to DB-25 convertor is required. This necessary cable is not shipped with all models that require it.
 - The convertor cable automatically ships with:
 - 7043-140
 - 7043-150
 - 7043-260
 - The convertor cable must be order separately for:
 - 7044-170
 - 7044-270
 - (# 3925) Serial Port Converter Cable 9-pin to 25-pin, pigtail

- (# 3926) Async Printer/Terminal Cable 9-pin to 25-pin, 4 m length

3.6.6 44P Model 270 Publications

Table 27 provides the publications available for the Model 270:

Table 27. Publications for the Model 270

Order Number	Title
SA38-0574	<i>7044 Model 270 Hardware Setup</i>
SA38-0573	<i>7044 Model 270 User's Guide</i>
SA23-2690	<i>Customer Support Information (US customers only)</i>
SA23-2652	<i>System Unit Safety Information</i>
SA38-0538	<i>PCI Adapter Placement Reference Guide</i>

Table 28 provides the additional publications that are available. To order, contact your IBM representative.

Table 28. Additional Publications for the Model 270

Order Number	Title
SA38-0572	<i>7044 Model 270 Service Guide</i>
SA38-0509	<i>Diagnostics Information for Multiple Bus Systems</i>
SA38-0516	<i>IBM RS/6000 Adapters, Devices and Cables Information for Multiple Bus Systems</i>

3.7 RS/6000 7046 Model B50 Overview

The following sections discuss the minimum configuration and optional features of the RS/6000 Model B50, which is shown in Figure 15.



Figure 15. Model 7046-B50

3.7.1 Model B50 Minimum Configuration

Table 29 lists the characteristics of a Model B50 minimum configuration.

Table 29. Model B50 Minimum Configuration

Model B50 Minimum Configuration	
Microprocessor	375 MHz PowerPC 604e
Level 1 (L1) cache	32 KB data/32 KB instruction
Level 2 (L2) cache	1 MB
RAM (minimum)	128 MB SDRAM
Memory bus width	64-bit
Integrated ports	Tablet, keyboard, mouse, 10/100 Ethernet, Ultra SCSI, serial (two), parallel, and stereo audio
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Four (two disk and two media bay available)
Expansion slots	Two PCI

Model B50 Minimum Configuration	
PCI bus width	32-bit
Memory slots	Four
CD-ROM drive	32X (Max) SCSI-2 CD-ROM
Diskette drive	1.44 MB 3.5-inch diskette drive
Service Processor	No
Operating Systems versions	4.3.3 unlimited user license standard
	Linux enabled
System dimensions and weight	3.5" H x 17.6" W x 24.1" D (standard 2U) (88 mm x 447 mm x 612 mm or 751.8mm (29.6") including covers, pull handle, cables and bend radius; 32 -35 lbs. (14.5-15.9 kg)
Warranty	24x7, on-site for one year (limited), at no additional cost

3.7.2 Model B50 System Expansion

Table 30 shows the possible storage and memory expansion configurations.

Table 30. Model B50 System Expansion

Model B50 System Expansion	
RAM	Up to 1 GB of SDRAM
Internal disk storage	Up to 54.6 GB by RPQ/special request
External disk storage	Options include: 2104 Expandable Storage Plus, IBM 2105 Enterprise Storage Server, IBM 7133 SSA Disk Subsystem. For additional information, see Chapter 13, "External Storage Architectures and Devices" on page 421.

3.7.3 Model B50 Optional Features

The following section outlines the optional features of the Model B50.

Optional features are defined as optional internal devices and internal adapters that can be configured on the RS/6000 Model B50.

Table 31 lists the optional internal features of the Model B50 and indicates whether each feature is available.

The status of a feature is indicative of these qualifications:

A Indicates features that are available and orderable on the specified models.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 31. Model B50 Optional Features

Feature Code	Description	Status
Memory		
4150	128 MB SDRAM DIMM	A
4169	256 MB SDRAM DIMM	A
Internal Disk Drives		
3025	9.1 GB 10K RPM Ultra2 SCSI Disk Drive Assembly	A
3026	18.2 GB 10K RPM Ultra2 SCSI Disk Drive Assembly	A
3112	9.1 GB Ultra SCSI 16-bit 1-inch (25 mm) High Disk Drive Assembly	A
3113	18.2 GB Ultra SCSI 16-bit 1-inch (25 mm) High Disk Drive Assembly	A
3129	36.4 GB Ultra SCSI 10K RPM 1 Inch	A
Internal CD-ROMs		
2624	CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
SCSI Adapters		
2494	Ultra2 SCSI 3-Channel RAID	A
6205	Dual Channel Ultra2 SCSI	A
6204	Ultra SCSI Differential (3.3V or 5.0V)	A

Feature Code	Description	Status
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential (older 5.0V)	A
SSA Adapters		
6231	128 MB DRAM Option Card (For 6225)	A
6235	32 MB Fast-Write Cache Option Card	A
6230	Advanced SerialRAID Plus Adapter	A
Asynchronous Adapters		
2943	8-Port Asynchronous Adapter EIA-232/RS-422	A
2944	128-Port Asynchronous Controller	A
ATM Adapters		
2963	Turboways 155 PCI UTP	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
4959	Token-ring adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	A
4951	4-Port IBM 10/100 Mbps Ethernet PCI Adapter	A
Digital Trunk Adapters		
2947	ARTIC960Hx 4-Port Selectable PCI Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A

Feature Code	Description	Status
Miscellaneous		
3752	Service Package (service publications, wrap plugs, and diagnostic CD-ROM)	A
6041	Mouse, 3 button, white	A
8741	Mouse, 3 button, black	A
6559	Additional Universal Mounting Rails (OEM/IBM Racks)	A

3.7.4 Model B50 Racks

The high-density RS/6000 Model B50 is designed to meet the height and width requirements of industry-standard racks for Internet service provider (ISP) and application service provider (ASP) companies. You can also install the B50 in an IBM 7014-S00 rack if you need additional frame depth.

3.7.4.1 Model B50 in S00 Rack

The Model S00 rack has 32 EIA units of vertical mounting space for 19-inch-wide drawers. Sixteen Model B50s can be stored into the S00 rack. 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. Ensure that you have enough outlets, and the PDU outlets are available in the rack to handle the configuration.

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.

The B50 may also be placed in the T00 or T42 racks from IBM.

3.7.4.2 B50 in Industry Standard Racks

In an industry-standard 2 meter 19-inch rack enclosure, you can package up to:

- Twenty servers (the number of servers that may be operated in a rack depends on individual customer environment).
- Sixty discrete network connections based on three 10/100 Ethernet connections for each Model B50 and 20 B50s in a rack.

Note

A software update is necessary if the 7046-B50 was installed prior to February 25, 2000 when using the 9513-AG1 monitor. To obtain the update access, go to:

<http://techsupport.services.ibm.com/rs6k/fixes.html>

3.7.5 Model B50 Configuration Notes

It is important to keep the following notes in mind when configuring the Model B50:

- PCI and graphics adapters are subject to placement rules, which are tabulated in the *Adapter Placement Guide* in this handbook.
- You can only have a maximum of 15 B50s in an S00 Rack since the extra PDU that will be required will take up one EIA. A maximum of 16 Model B50s will physically fit into the S00 rack, but all cannot be plugged into the PDUs of the rack.
- Up to four memory DIMMs (128 MB, 256 MB) can be installed. These DIMMs do not have to be installed in pairs.

3.7.6 Model B50 Publications

Table 32 provides the publications that for the Model B50.

Table 32. Publications for the Model B50

Order Number	Title
SA38-0562	<i>RS/6000 Enterprise Server B50 Setup Guide</i>
SA38-0563	<i>RS/6000 Enterprise Server B50 User's Guide</i>
SA23-2690 (USA customers only)	<i>Customer Support Information</i>
SA23-2652	<i>System Unit Safety Information</i>
SA38-0538	<i>PCI Adapter Placement Reference Guide</i>

These additional publications listed in Table 33 can be ordered from IBM.

Table 33. Additional Publications for the Model B50

Order Number	Title
G221-7099	<i>RS/6000 - Enterprise Server Model B50 Spec Sheet</i>

Order Number	Title
SA38-0564	<i>RS/6000 7046 B50 Service Guide</i>
SA38-0550	<i>7014 Model S00 Rack Installation and Service Guide</i>

Chapter 4. Midrange Enterprise Servers

This chapter takes a look at the following models and solutions:

- The Deskside Servers:
 - 7025-F50 - Enterprise Server
 - 7025-F80 - Enterprise Server
- The Rack-Mounted Servers:
 - 7026-H50 - Enterprise Server
 - 7026-H70 - 64-Bit Enterprise Server
 - 7026-HA-H70 - Enterprise Server Solution
 - 7026-H80 - 64-Bit Enterprise Server
 - 7026-HA-H80 - Enterprise Server Solution
 - 7026-M80 - 64-Bit Enterprise Server
 - 7026-HA-M80 - Enterprise Server Solution

The Midrange Enterprise Server can be defined as a system with an ability to expand the numbers of CPUs and storage capabilities to service a medium- to large-size company or department.

4.1 Model History and Uses

This section describes the deskside models of RS/6000 in the Enterprise Server group in addition to the midrange rack-mounted servers in the Enterprise Server group. The intention of outlining them is to present their development from a historical perspective.

4.1.1 The Deskside Servers

The following sections describe the servers that are designed to fit comfortably inside your office environment.

4.1.1.1 Model F50

The Model F50 became available April 25, 1997 and even further expanded the machine's capabilities with the 4-way SMP and faster processors with internal Serial Storage Architecture (SSA). This commercial solution addresses mission-critical, online transaction processing (OLTP) and collaborative computing comprised of Lotus Notes, Domino Server, Internet, intranet, extranet, and groupware application solutions.

4.1.1.2 Model F80

The Model F80 became available June 09, 2000 and is a major update in performance over the Model F50 it is designed to replace. The Model F80 allows you to grow to a 6-way SMP. With redundant power and cooling, 64-bit architecture, and internal storage options, the F80 provides the needed capacity for growth.

4.1.2 The Rack-Mounted Servers

The RS/6000 7026 Model H10 (withdrawn from marketing on February 27, 1998) was the first entry workgroup server in a 7 EIA drawer that was made available on February 14, 1997. The Model H10 came in one-way or two-way PowerPC 604e configurations running at 166 MHz. The service processor was standard on all rack-mounted servers since the introduction of the H10.

4.1.2.1 Model H50

The RS/6000 Enterprise Server 7026 Model H50 was announced on February 09, 1998. The Model H50 can be used as a stand-alone, multiuser, application, or database server and has the connectivity to participate in most currently installed UNIX and PC networks. It can also be configured for use as an SCSI, Ultra SCSI, or Serial Storage Architecture (SSA) RAID server by taking advantage of the under the covers storage expansion capability, which allows for system growth. The H50 provides enhanced reliability with its redundant cooling design and optional redundant power when ordered with a second internal power supply.

The H50 has been referred to as the rack version of the F50. However, there are some fundamental differences, such as there are three 6-packs of disks in the F50 and two 6-packs in the H50. Also, the H50 has redundant cooling and optional redundant power to provide high availability for critical application systems. The H50 can be inserted into any existing RS/6000 rack.

4.1.2.2 Model H70

Since its announcement April 6, 1999, the 64-bit H70 is an improved follow-on machine to the popular Enterprise Server Model H50. Growing business requirements can now be satisfied with enhanced 64-bit processors (RS64 II) introduced for the first time in a mid-size RS/6000 SMP server.

The Model H70 system, with one to four 64-bit 340 MHz processors and up to 8 GB of SDRAM memory, provides significant performance gains when coupled with 64-bit database and application programs or when running legacy 32-bit applications.

The H70 incorporates additional enhancements over the H50, including a faster I/O subsystem, integrated 10/100 Ethernet and Ultra SCSI Adapters and an additional L2 cache.

For those applications that are I/O performance intensive, the H70 offers improvement in the I/O data bandwidth of over 30 percent in comparison to the H50, providing 528 MB/s peak performance for those critical interactive and transaction-oriented applications, including Web-based access to databases for e-business transactions. The four PCI controllers, operating together on a multiple peer 64-bit bus at 66 MHz, provide this result.

The H70 is an eight EIA rack drawer installable in any existing RS/6000 rack. It has a medium-cost level of price performance. Redundant cooling and redundant power are designed into this system.

4.1.2.3 HA-H70 Cluster Server

The RS/6000 HA-H70 is a special package-priced, high-availability solution, with options for applications, services, and financing. The HA-H70 is not a new server model. The HA-H70 consists of currently available IBM products.

Some of its features are:

- Complete high-availability RS/6000 hardware and software solution platforms for mission-critical and e-business applications
- Package-priced RS/6000 high-availability solution, less than the sum of the parts
- Estimated 99.999 percent availability, or less than six minutes of annual unplanned system hardware and software downtime, estimated 99.998 percent availability, or less than eleven minutes of annual unplanned system and DB2 database downtime

4.1.2.4 Model H80

The IBM RS/6000 Model H80 is a member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM available June 09, 2000. A follow-on to the popular Model H70, the H80 provides the power, capacity, and expandability required for e-business mission-critical computing. The H80 can assist you in managing the evolution of your business to incorporate the power of the Web and 64-bit applications into your computing environment while still supporting existing 32-bit applications.

The H80 is designed to be a cost-efficient growth path to the future. It provides 64-bit scalability using the 64-bit RS64 III processor packaged in 1-, 2-, 4-, and 6-way processor configurations and incorporates an I/O

subsystem supporting 32-bit and 64-bit standard PCI adapters. The H80 has 1-, 2-, and 4-way processors that operate at 450 MHz while the 6-way processor operates at 500 MHz. The 1-way processor provides 2 MB of L2 cache with the 2-, 4-, and 6-way processors incorporating 4 MB of L2 cache per processor.

The H80 is packaged as a rack-mounted Central Electronics Complex (CEC) drawer cable attached to rack-mounted Remote I/O (RIO) drawers. The CEC and I/O drawers offer redundant power options and redundant cooling. The CEC drawer incorporates the system processors, memory, and supporting systems logic.

The Model H80 CEC and I/O drawers are mounted in the IBM 7014 rack. Additional space in the 7014 rack may be utilized for various storage and communications subsystems.

4.1.2.5 Model M80

The RS/6000 Model M80 is a new, mid-range member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM made available June 09, 2000. Positioned between the new Model H80 and the powerful Model S80, the M80 provides the power, capacity, and expandability required for e-business mission-critical computing. The M80 can assist you in managing the evolution of your business to incorporate the power of the Web and 64-bit computing into your environment while still supporting existing 32-bit applications.

The M80 delivers a cost-efficient growth path to the future. It provides 64-bit scalability via the 64-bit RS64 III processor packaged as 2-way and 4-way cards. With its two processor positions, the M80 can be configured into 2-, 4-, 6-, or 8-way SMP configurations. The M80 also incorporates an I/O subsystem supporting 32-bit and 64-bit standard PCI adapters. The M80 has 2- and 4-way processor cards that operate at 500 MHz and incorporate 4 MB of L2 cache per processor.

The M80 is packaged as a rack-mounted Central Electronics Complex (CEC) drawer cable attached to rack-mounted Remote I/O (RIO) drawers. The CEC and I/O drawers offer redundant power and redundant cooling. The CEC drawer incorporates the system processors, memory, and supporting systems logic.

The M80 CEC and I/O drawers are mounted in the IBM 7014 rack. Additional space in the 7014 rack may be utilized for various storage and communications subsystems.

4.1.3 Upgrade Paths

The model upgrade paths of the 7025 machines are as follows:

- F40 -> F50
- F50 -> H70
- F50 -> F80

The model upgrade paths of the 7026 machines are as follows:

- H50 -> H70
- H70 -> H80
- H80 -> M80

Keep in mind that there are upgrades within each model as it evolved to faster processors or if the client wanted to add more processors to enhance the machine's capabilities. However, these do not constitute model upgrades.

4.2 RS/6000 7025 Model F50 Overview

This section describes the standard features of the base Model F50 (Figure 16) in addition to the supported optional features.



Figure 16. Model 7025 F50 - Front View

4.2.1 Model F50 Standard Features

This section describes the features of a Model F50 standard configuration.

Table 34 provides the standard features of a Model F50.

Table 34. Model F50 Standard Configuration

Model F50 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 (minimum)	128 MB ECC synchronous DRAM
Memory bus width	128-bit

Model F50 Standard Configuration	
Standard Ports	One parallel, three serial, keyboard, mouse, Ethernet 10Base5/T
Internal disk drive	9.1 GB Ultra SCSI
Disk / Media bays	18 (one used)/four (two used)
Expansion slots	Nine (seven PCI, two PCI/ISA)
PCI bus width	Two 32-bit and one 64-bit
Memory slots	Two
CD-ROM drive	20X (Max) SCSI-2
Service Processor	Yes
Diskette drive	1.44 MB 3.5" diskette drive
SCSI Adapters	Two integrated SCSI-2 F/W adapters
AIX operating system version	Version 4.3 unlimited user license.
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) Weight will vary based on installed options
Warranty	24 x 7, on-site for one year (limited) at no additional cost

Figure 17 on page 76 shows the ports and slots that are visible from the rear view of the Model F50.

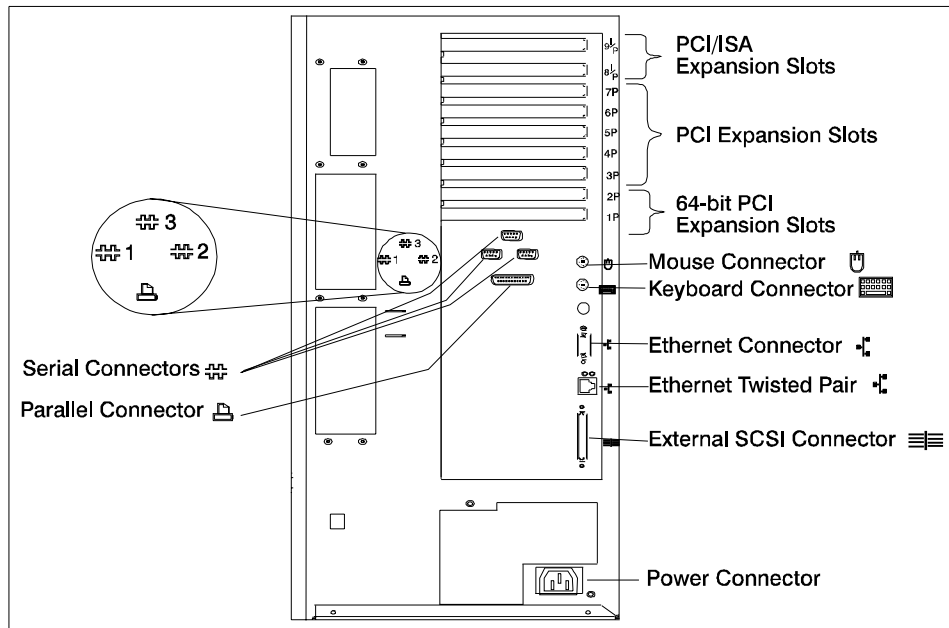


Figure 17. Model 7025 F50 - Rear View

4.2.1.1 Model F50 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 that users can take disks off-line 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 provide 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.

- When Service Director for RS/6000 is implemented (available at no additional charge if your RS/6000 processor is covered by an IBM Warranty or IBM Maintenance Service Agreement) it can monitor and analyze system errors and, if needed, 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.

4.2.1.2 Model F50 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, to call the IBM Service Center. The objective is to predict a potential component failure and automatically dispatch IBM Service to take preventive 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.

4.2.2 Model F50 System Expansion

Table 35 shows the possible maximum processor, storage, and memory configurations.

Table 35. Model F50 System Expansion

Model F50 System Expansion	
SMP configurations	Two, three, or four -way CPU cards
RAM	Up to 3 GB
Internal disk storage	Up to 336.7 GB (327.6 hot-swappable) per system

4.2.3 Model F50 Optional Features

This section describes the internal optional features that can be added to the Model F50 base configurations at an additional cost.

The status of a feature is indicative of these qualifications:

- A** To indicate features that are available and orderable on the specified models.
- S** Indicates a supported device in this model.
- W** Indicates a withdrawn feature. This feature was previously available.
- R** Indicates this adapter is not supported and must be removed during a model conversion from an F40.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards and cables, are not included.

Table 36 provides the potential optional features available on the Model F50.

Table 36. Model F50 Optional Features

Feature Code	Description	Status
Processors		
4358	1-Way 332 MHz 604e Processor	A
4359	2-Way 332 MHz 604e Processor	A
4356	1-Way 332 MHz 604e Processor (Select)	A
4357	2-Way 332 MHz 604e Processor (Select)	A
1001	Support Processor	R
4341/4340	1-Way 233 MHz 604e Processor	R
4306/9305	1-Way 166 MHz 604e Processor	R
Memory		
4111	8 MB DRAM DIMM	R
9082	16 MB DRAM DIMM (Base)	R
9080	128 MB DRAM DIMM	R
9083	128 MB (4 x 32 MB) SDRAM DIMM	A
4112	16 MB DRAM DIMM	R
4113	32 MB DRAM DIMM	R
4128	32 MB DRAM DIMM (Select)	R

Feature Code	Description	Status
4114	64 MB DRAM DIMM	R
4129	64 MB DRAM DIMM (Select)	R
4107	64 MB (2 x 32 MB) SDRAM DIMM	A
4115	128 MB DRAM DIMM	R
4130	128 MB DRAM DIMM (Select)	R
4110	256 MB (2 x 128 MB) SDRAM DIMM	A
4106	256 MB (2 x 128 MB) SDRAM DIMM	A
4093	Memory Board, 16 Position Level 3.0	A
Internal Disk Bays		
6520/6537	SCSI 6-pack hot swap	R
6523	SCSI 6-pack non-hot swap	R
6519	SCSI 6-pack hot swap	A
6538	SSA 6-pack select	A
6527	SSA 6-pack select	S
6539	SSA 6-pack	A
Internal Disk Drives		
3083/9398	2.2 GB Hot-Swap F/W	S
3084/3092	4.5 GB Hot-Swap F/W	S
3080/3081	4.5 GB 1" Hot-Swap F/W	S
3090/2912	9.1 GB Hot-Swap F/W	S
2901	4.5 GB Ultra SCSI Hot-Swap	W
2913/2919	9.1 GB 1" Ultra SCSI Hot-Swap	A
2908	9.1 GB 1" Ultra SCSI Hot-Swap	A
2916	9.1 GB 1" Ultra SCSI Hot-Swap	S
2918	18.2 GB 1" Ultra SCSI Hot-Swap	S
3101	18.2 GB 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A

Feature Code	Description	Status
3071	4.5 GB SSA Hot-Swap	W
3070	9.1 GB 10K RPM SSA Hot-Swap	A
3074	9.1 GB 1" SSA Hot-Swap	W
Internal Tape Drives		
6137	24/48 GB 4 mm Autoloader	R
2630	1/4" Tape 1.2 GB SCSI	R
6154	20/40 GB 8 mm (White)	A
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2619	20X Speed CD-ROM	A
Graphics Accelerators		
2838	GXT120P	W
2830	GXT130P	A
2853	GXT800P	W
2859	GXT800P with Texture Graphics	W
2851	GXT250P	R
2852	GXT255P	R
2854	GXT500P	R
2855	GXT550P	R
2837	MVP Multi-Monitor	R
2856	7250-001/002 GXT1000	R
8679	Multimedia Kit	R
2638	Video Capture	R
SCSI Adapters		
6206	Ultra SCSI	A
6207	Ultra SCSI Differential	A
2493	SCSI 2 Fast / Wide RAID	W

Feature Code	Description	Status
6205	Ultra2 SCSI PCI	A
2494	Ultra SCSI 3-Channel RAID	A
SSA Adapters		
6222	SSA Fast-Write Cache Option	A
6235	32 MB Fast-Write Cache Option Card	A
6231	128 MB Cache Option	A
6230	SSA 2-Way RAID 5	A
6232	SSA Boot Enable Feature	A
Fibre Channel Adapter		
6227	PCI Fibre Channel Adapter	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
System Adapter		
2751	S/390 ESCON Channel PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
2998	Turboways 25 PCI	A
Token-ring adapters		
4959	Token-ring adapter	A
Ethernet Adapters		
2968	10/100 Mbps	A

Feature Code	Description	Status
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2987	Ethernet AUI / RJ-45	A
4951	4-Port Ethernet Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
Miscellaneous		
6308	Personal Dictation Microphone	R
8741	Mouse, 3 button, black	A
6041	Mouse, 3 button, white	A
3752	Service Publications	A
6405-6409	Powered Speakers	R
6549	Additional Power for 2nd and 3rd 6-packs	R

4.2.3.1 Model F50 Additional Processors

An additional 1-way or 2-way 332 MHz processor can be added. All processors must have the same clock rate; they cannot be mixed.

4.2.3.2 Model F50 Additional Memory

The base configuration can consist either of a 128 MB (4-32 MB 200-pin SDRAM DIMMs) or 256 MB (2 x 128 MB) Error Checking and Correcting (ECC) SDRAM DIMMs. Both 10 ns memory. 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.

4.2.3.3 Model F50 Optional Disk Bays

The F50 features four media bays, six standard, and twelve optional, half-height, 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 SSA 6-pack (# 6539) or SCSI hot swap 6-pack (# 6519), depending on which optional bays are going to be installed. In addition, one connection cable for each 6-pack must be ordered. Feature # 2444, 2447, or 8447 provide connection to SCSI implementations. Feature # 2543 or 2450 provide cable connections for SSA implementations.

Figure 18 shows the internal disk and media bays available on the Model F50.




Figure 18. Model 7025 F50 - Internal Disk and Media Bays

4.2.4 Model F50 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.

- There can be a maximum of two graphics accelerators 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.

4.2.5 Model F50 Publications

Table 37 provides the publications for the Model F50:

Table 37. Publications for the Model F50

Order Number	Title
SA38-0539	<i>IBM RS/6000 7025 Setup Procedure</i>
SA38-0540	<i>IBM RS/6000 7025 User's Guide</i>
SA23-2652	<i>System Unit Safety Information</i>
SA38-0538	<i>PCI Adapter Placement Reference Guide</i>

4.3 RS/6000 7025 Model F80 Overview

This section describes the standard features of the base Model F80 (Figure 19) in addition to the supported optional features.



Figure 19. Model 7025 F80 - Front View

4.3.1 Model F80 Standard Features

This section describes the features of a Model F80 standard configuration.

Table 38 provides the minimum configuration of a Model F80.

Table 38. Model F80 Minimum Configuration

Model F80 Minimum Configuration	
Microprocessor	450 MHz (1-, 2-, or 4-way) or 500 MHz (6-way) RS64 III
Level 1 (L1) cache	128 KB data/128 KB instruction
Level 2 (L2) cache	2 MB (1-way), 4 MB (2-, 4-, or 6-way) ECC

Model F80 Minimum Configuration	
RAM (minimum)	256 MB ECC synchronous DRAM
Memory bus width	128-bit
Standard ports	One parallel, four serial, keyboard, mouse, Ethernet 10/100BaseT
Internal disk drive	9.1 GB Ultra2 SCSI or SSA disk drive
Disk / Media bays	12 ¹ / 3 (two used)
Expansion slots (hot-plug)	10 (6 3.3V at 66 MHz, 4 5V at 33 MHz)
PCI bus width	All 64-bit
Memory slots	Two
CD-ROM drive	32X (Max) SCSI-2
Service processor	Yes
Diskette drive	1.44 MB 3.5" diskette drive
SCSI adapters	SCSI-2 F/W (internal) Ultra2 SCSI (external VHDCI)
AIX operating system version	Version 4.3.3 ² unlimited user license
System dimensions	24.0" H x 19.0" W x 28.7" D (610 mm x 483 mm x 728 mm); 65 lbs (30 kg) Weight will vary based on installed options
Warranty	On-site for one year (limited) at no additional cost
¹ Two six-packs (two additional bays with Dual Bay Mounting Kit) ² AIX 4.3.3 and the APARs IY09047 and IY09814	

Figure 20 on page 87 shows the ports and slots that are visible from the rear view of the Model F80.




Figure 20. Model 7025 F80 - Rear View

4.3.1.1 Model F80 Processor Subsystem

The Model F80 allows a 1-, 2-, 4-, or 6-way configuration. For 1-, 2- and 4-way configurations, the processor cards are equipped with 450 MHz RS64 III processors. The 6-way configuration comes with 500 MHz RS64 III. Both variants of RS64 III processors have 128 KB of both instruction and data Level 1 (L1) cache. Concerning the L2 cache, all configurations, except the 1-way, have 4 MB of L2 cache (4-way set associative). The 1-way machine is only equipped with 2 MB L2 cache (4-way set associative).

4.3.1.2 Model F80 Memory

The base configuration consists of a 256 MB (on the 1-way CPU-board or the memory board) of Error Checking and Correcting (ECC) SDRAM DIMMs in a 1-way configuration. If there is more than one processor in the system, the memory must be installed in quads on a memory board; so, the base configuration for these systems consists of 512 MB (4 x 128 MB). A maximum of 16 GB of memory is supported when both memory boards are used.

Note

A one-way system is limited to only one memory card, which can only be used when the one-way processor card contains no DIMMs.

4.3.1.3 Model F80 Reliability, Availability, Serviceability Features

The Model F80 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 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 slots in the Model F80 are hot-plug PCI I/O slots bringing a new and increased level of availability for this system. The need to power down and reboot is eliminated when adding or replacing most PCI adapters, thus reducing time delays while keeping network and application processing up and running.
- The system, I/O, and PCI buses are designed with parity error detection.
- The disk drives are hot-swappable so that users can take disks off-line or put them back into service without interrupting the operation of their system (except the disks in Dual Bay Mounting Kit, if used).
- RAID 0, 1, and 5 for internal devices (SCSI or SSA), which provides data integrity and fault tolerance in the event of disk drive failure.
- Online (concurrent) and remote (LAN or modem) system diagnostics.
- Dynamic Processor Deallocation
- 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.
- Hot-plug redundant power and cooling capability. The power supplies and fans can be replaced while the system is running when the additional redundant power supply is installed.
- When Service Director for RS/6000 is implemented (available at no additional charge if your RS/6000 processor is covered by an IBM Warranty or IBM Maintenance Service Agreement), it can monitor and analyze system errors and, if needed, 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.

4.3.1.4 Model F80 Enhanced 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, call IBM Service Center capability. The system objective is to predict a potential component failure and automatically dispatch an IBM service representative to take preventive maintenance measures to avoid an outage. The system can also call the IBM Service Center in the event of a critical failure.
- Integrated system monitoring enhances failure error logging, analysis, and alert.
- Can be programmed by the system administrator to re-boot, in most instances, after experiencing the following conditions:
 - AC power restored (after a power loss)
 - Hardware checkstop failures
 - Machine check interrupt
 - Operating system hang or failure

The Model F80's enhanced service processor is functional in standby mode and does not require a functional system for operation. This function provides enhanced RAS by not requiring AIX to be operational for interfacing with a remote system administrator or Service Director for RS/6000.

4.3.2 Model F80 Optional Features

This section describes the internal optional features that can be added to the Model F80 base configurations at an additional cost.

The status of a feature is indicative of these qualifications:

- A** To indicate features that are available and orderable on the specified models.
- S** Indicates a supported device in this Model
- R** Indicates this adapter is not supported and must be removed during a model conversion from an F50.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards and cables, are not included.

Table 39 provides the potential optional features available on the Model F80.

Table 39. Model F80 Optional Features

Feature Code	Description	Status
Processors		
5201	1-way 450 MHz RS64 III Processor	A
5202	2-way 450 MHz RS64 III Processor	A
5204	4-way 450 MHz RS64 III Processor	A
5206	6-way 500 MHz RS64 III Processor	A
Memory		
4075	Memory Board	A
4093	Memory Board, 16 Position Level 3.0	R
4107	64 MB (2 x 32 MB) SDRAM DIMM	A
9083	128 MB (4 x 32 MB) SDRAM DIMM	S
4110	256 MB (2 x 128 MB) SDRAM DIMM	A
4106	256 MB (2 x 128 MB) SDRAM DIMM	S
4119	512 MB (2 x 256 MB) SDRAM DIMM	A
4131/4100	1024 MB (2 x 512 MB) SDRAM DIMM	A
Internal Disk Bays		
6555	SCSI Dual IPL Bay	A
6553	SCSI 6-pack hot swap	A
6519	SCSI 6-pack hot swap	R
6554	SSA 6-pack	A
6527/6538	SSA 6-pack select	R
6539	SSA 6-pack	R
Internal Disk Drives		
3083/9398	2.2 GB Hot-Swap F/W	R

Feature Code	Description	Status
2900/9394	4.5 GB Ultra SCSI	S
3084/3092	4.5 GB Hot-Swap F/W	R
3080/3081	4.5 GB Hot-Swap F/W	S
3005	4.5 GB Ultra SCSI	S
2901	4.5 GB Ultra SCSI Hot-Swap	S
2902	4.5 GB Ultra SCSI Hot-Swap	S
3006	4.5 GB Ultra SCSI F/W	S
3071	4.5 GB SSA Hot-Swap	S
3090/2912	9.1 GB Hot-Swap F/W	S
3091	9.1 GB Hot-Swap F/W	S
3116	9.1 GB Ultra SCSI	S
3115	9.1 GB Ultra SCSI	S
2908	9.1 GB Ultra SCSI	S
2911/3019	9.1 GB Ultra SCSI Hot-Swap	S
2913/2919	9.1 GB Ultra SCSI Hot-Swap	S
2916	9.1 GB Ultra SCSI Hot-Swap	S
3002	9.1 GB Ultra SCSI 10K RPM	S
3025	9.1 GB Ultra SCSI 10K RPM	A
3008	9.1 GB Ultra SCSI Hot-Swap 10K RPM	S
3021	9.1 GB Ultra2 SCSI	A
3072	9.1 GB SSA Hot-Swap	R
3074	9.1 GB SSA Hot-Swap	S
3079	9.1 GB SSA 10K RPM	A
3070	9.1 GB SSA Hot-Swap 10K RPM	S
3078	9.1 GB SSA Hot-Swap 10K RPM	S
2918	18.2 GB Ultra SCSI Hot Swap	S
3101	18.2 GB Ultra SCSI Hot Swap	S

Feature Code	Description	Status
3104	18.2 GB Ultra SCSI Hot Swap	S
3023	18.2 GB Ultra2 SCSI	A
3026	18.2 GB Ultra2 SCSI 10K RPM	A
3077	18.2 GB SSA 10K RPM	A
3129	36.4 GB Ultra2 SCSI 10K RPM	A
Internal Tape Drives		
6142	4/8 GB 4 mm	S
6159	12/24 GB 4 mm	A
6147	5/10 GB 8 mm	S
6154	20/40 GB 8 mm (White)	S
6156	20/40 GB 8 mm (Black)	A
Internal CD-ROMs		
2618	8X Speed CD-ROM	S
2619	20X Speed CD-ROM	S
2624	32X/40X Speed CD-ROM	A
Graphics Accelerators		
2839	GXT110P	R
2838	GXT120P	S
2830	GXT130P	A
2853	GXT800P	R
2859	GXT800P with Texture Graphics	R
SCSI Adapters		
6204	Ultra SCSI Differential	A
6205	Ultra2 Dual Channel SCSI	A
6206	Ultra SCSI	A
6207	Ultra SCSI Differential	S
6208	SCSI 2 Fast / Wide	S

Feature Code	Description	Status
6209	SCSI 2 Fast / Wide Differential	S
2493	SCSI 2 Fast / Wide RAID	S
2494	Ultra SCSI 3-Channel RAID	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	S
6218	SSA 4-Port RAID Adapter	R
6222	SSA Fast-Write Cache Option	A
6235	32 MB Fast-Write Cache Option Card	A
6225	IBM Advanced SerialRAID	S
6230	IBM Advanced SerialRAID Plus	A
6231	128 MB DRAM Option Card	A
Fibre Channel Adapters		
6227	Fibre Channel Adapter	A
Async Adapters		
2931	8-Port Async EIA-232 (ISA)	R
2932	8-Port Async EIA-232/422 (ISA)	R
2933	128-Port Async Controller (ISA)	R
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
2948	ARTIC960Hx 4-Port T1/E1 PCI	R
2949	ARTIC960Hx DSP Resource	R
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
6309	Digital Trunk Quad Adapter	R
6311	ARTIC960RXF Digital Trunk Adapter	A

Feature Code	Description	Status
HIPPI Adapters		
2732	Serial Hippi Long-Wave Adapter	R
2733	Serial Hippi Short-Wave Adapter	R
System Adapter		
2751	S/390 ESCON Channel PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
2998	Turboways 25 PCI	R
Token-ring adapters		
2920	4/16 token-ring adapters	S
2979	Auto LANStreamer token ring	S
4959	4/16 token-ring adapters	A
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2986	Fast Etherlink XL 3Com	S
2987	Ethernet AUI / RJ-45	A
4951	4-Port 10/100 Mbps	A
WAN Adapters		
2701	4-Port Communications Controller (ISA)	R
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A

Feature Code	Description	Status
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
X.25 Adapter		
2961	X.25 Adapter (ISA)	R
Miscellaneous		
6548	Redundant AC Power Supply	A
8741	Mouse, 3 button, black	A
6041	Mouse, 3 button, white	A
3752	Service Publications	A

Note

Disks # 3090, # 3091, and # 2912 are 1.6 inch high and need two disk bays.

4.3.2.1 Model F80 Additional Processors

Customers with an installed Model F80 may increase the number of processors in their system up to a maximum of a 6-way SMP using feature conversions. Such conversions are accomplished by replacing the system's existing processor card with a processor card containing a greater number of processors. The existing Model F80 processor card being replaced is returned to IBM.

4.3.2.2 Model F80 Additional Memory

The Model F80 has two possible memory boards with 16 DIMM slots each. The memory boards must be filled in units of four (quads) with the same size DIMM. Either 512 MB (4 x 128 MB # 4110), 1024 MB (4 x 256 MB # 4119), or 2048 (4 x 512 MB # 4131) SDRAM DIMMs (200 pin, 10 ns) can be added. When upgrading from an F50, it is also possible to use 32 MB DIMMs (# 4107). Memory feature # 4107 is available only on 7026-H50 to 7026-F80 model conversions to complete the F80 requirements for memory to be installed in quads when installed on a 16-Position Memory boards feature.

For a 1-way configuration, it is also possible to use the eight DIMM slots on the 1-way processor card. These slots can be filled in pairs instead of quads.

If DIMMs are installed on the processor card, it is not possible to install DIMMs on the optional memory boards.

4.3.2.3 Model F80 Disk Bays

The Model F80 features three media bays, twelve hot-swap disk bays, and an additional boot disk bay. Two of the three media bays are used for the diskette drive and CD-ROM drive. The third one can be used for an additional drive, such as a tape drive. The twelve disk bays are divided into two 6-packs, which can be either SCSI (# 6553) or SSA (# 6554). SCSI and SSA 6-packs can be mixed. The two additional boot disk bays (SCSI only) can be configured for boot disk mirroring. These two disks do not impact the two 6-packs; so, the 6-packs can be used, for example, as RAID. In addition, one connection cable for each 6-pack and the boot disk bay (if mirroring the boot disk) must be ordered. Feature # 2421, # 2432, and # 3109 provide connection to SCSI implementations. Feature # 2454 provides cable connections for SSA implementations.

Figure 21 shows the internal disk and media bays available on the Model F80.

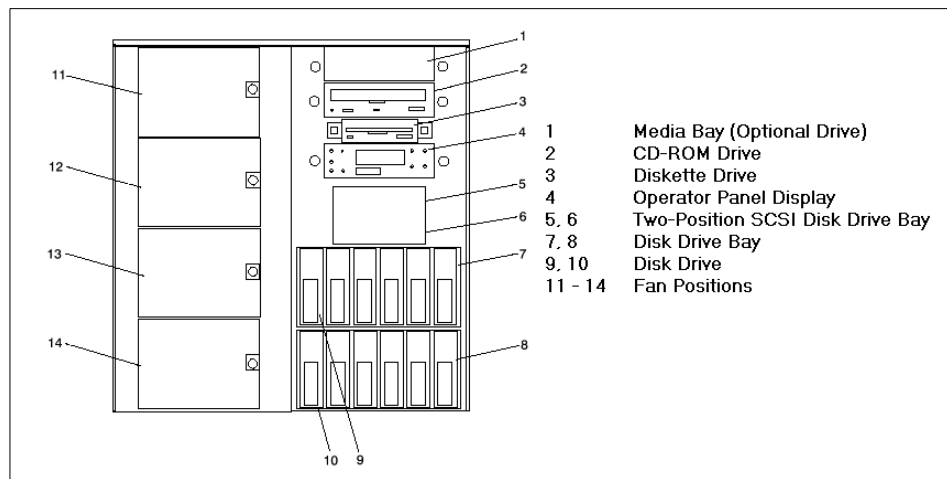


Figure 21. Model 7025 F80 - Internal Disk and Media Bays

4.3.2.4 Model F80 Internal Disk Drives

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 9.1 GB, 18.2 GB, or 36.4 GB capacities. Some older disks can also be used in the Model F80 (refer to Table 36 on page 78).

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 disks occupy two. For this reason, the maximum number of full-height disk drives that can be installed is seven (two bays for each drive), and the maximum number of half-height disk drives is 14 (one bay for each drive). Since all the 9.1 GB and 18.2 GB SSA disks are half height, these disks can each occupy a single disk bay in the 6-packs, thus making a total of 12 disks (the additional bootdisk can only be a SCSI disk).

4.3.2.5 Model F80 Internal Tape Drives

Four types (4 GB/8 GB 4 mm (DDS-2), 12 GB/24 GB 4 mm (DDS-3), 5 GB/10 GB 8 mm, or 20 GB/40 GB 8 mm) of internal tape drives can be installed in the free internal media bay. The 4 GB/8 GB 4 mm (DDS-2) and the 5 GB/10 GB 8 mm have been withdrawn from marketing.

4.3.2.6 Model F80 Serial Storage Architecture (SSA) and SSA RAID

The Model F80 supports both external and under-the-covers SSA disk subsystems. The SSA Advanced SerialRAID Plus Adapter (# 6230) can be used to control both 6-packs of SSA disks. The Feature # 6230 supports boot from SSA but only in non-RAID configurations. When implementing SSA internally, SSA hard disks (# 3025, # 3021, # 3026, and # 3023) and 6-pack backplane (# 6554) are required. SSA disks and backplanes are hot-swap capable.

4.3.2.7 Model F80 Graphics Accelerators

For 2D graphics applications, the Model F80 offers the POWER GXT130P Graphics Accelerator. The POWER GXT130P is a versatile, low-priced graphics accelerator. It is a half-length PCI card with a standard 15-pin DIN output connector and can support multisync monitors with at least a 64 KHz horizontal scan capability or DDC-2B-compliant multisync monitors with at least a 38 KHz horizontal scan capability. For Model F80, it is a cost-effective choice for server environments that require a graphical user console.

The old GXT120P adapter is also supported for use, but this adapter is withdrawn from marketing and, therefore, not orderable.

Note

No multimedia features (audio or video) are supported in the Model F80.

4.3.2.8 Model F80 Redundant Power Supply and Cooling

A Model F80 standard configuration contains two power supplies and two cooling fans to support up to 6-way SMP systems with 16 GB of memory. An optional third power supply with two additional cooling fans (# 6548) can be added to provide redundant power and cooling. All power supplies and fans are hot-pluggable and allow replacement while the system is running when the additional redundant power supply is installed.

4.3.3 Model F80 Configuration Notes

When planning or implementing an RS/6000 Model F80, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

- Old 1.6-inch disks (when supported) require two hot-swappable bays.
- There can be a maximum of only one graphics accelerators installed in the system.
- A SCSI boot disk must be configured in the additional boot disk bay if both 6-packs are used in RAID mode (SCSI or SSA) and if you do not want to boot from external storage.

4.3.4 Model F50 Upgrades

Current owners of RS/6000 7025 Model F50 may order an upgrade to convert their system to Model F80 while keeping the original system serial number. The upgrade to Model F80 systems will require replacement of all system processors. The first Model F50 processor card can be replaced with one Model F80 2-way or greater processor card using a feature conversion. The memory DIMMs and most of the adapters can also move from Model F50 to Model F80. The following points should be considered:

- Memory DIMMs in the Model F80 system must be installed in units of four (quads). Only in a one-way configuration is it possible to install up to four pairs of DIMMs on the one-way processor card.
- Most internal disks can also move to Model F80, but they need a new carrier.
- Only the Graphics Accelerators GXT120P and GXT130P are supported.
- AIX 4.3.3 with APARs IY09047 and IY09814 is required. Old system disks might be upgraded before use in Model F80.

4.3.5 Model F80 Publications

Table 40 provides the publications for the Model F80:

Table 40. Publications for the Model F80

Order Number	Title
SA38-0569	<i>RS/6000 Enterprise Server Model F80 Installation Guide</i>
SA23-2652	<i>System Unit Safety Information</i>

Table 41 provides the additional publications for the Model F80.

Table 41. Additional Publications for the Model F80

Order Number	Title
SA38-0567	<i>RS/6000 Enterprise Server Model F80 User's guide</i>
SA38-0568	<i>RS/6000 Enterprise Server Model F80 Service guide</i>

The publications previously listed are also available for download from the Internet. If you do not know where to get the publications, ask your sales representative.

4.4 High-Availability Cluster Server Solution HA-F80

This section describes the hardware and software that completes the HA-F80 solution.

The HA-F80 is a two-machine, high-availability cluster consisting of two RS/6000 Model F80s with AIX Version 4.3.3 or later system software and HACMP 4.3 or later high-availability cluster software. One HA-F80 configuration includes the Model 7133-T40 SSA Disk Storage Subsystem. Additional storage options can be configured with the HA-F80.

The HA-F80 solution package must include a minimum of the following hardware and software that must be ordered together as specified to qualify for the HA-F80 package price:

- Two 7025-F80 Enterprise Servers, each incorporating:
 - One HA Solution Indicator (# 0700)
 - One HA Solution Processor Card - Choose from:
 - One HA Solution Processor Card, 4-Way (# 0505)
 - One HA Solution Processor Card, 6-Way (# 0508)
- 512 MB Memory

- Two PCI SSA Adapters
- Two LAN Adapters
- One Redundant Power Supply (# 6548)
- One Serial to Serial Port Cable (# 3125)
Only one serial cable (# 3125) is needed for the two F80 servers.
Redundant dedicated heartbeat and messaging paths are required via LAN and ASYNCH ports on each server.
- One Software Preinstall (# 5005)
Redundant SCSI or SSA IPL disks are required. AIX and HACMP software will be preinstalled on 7133 SSA disks unless optional SCSI disks are ordered.
- One AIX Version 4.3.3, or later license
- One HACMP Version 4.3, or later license

Note

Additional optional 7025-F80 features can be added.

- One 7133-T40 Serial Disk System including:
 - Four SSA Disk Drives
 - Six Advanced SSA Cables

Note

Additional optional 7133-T40 features can be added.

4.5 RS/6000 7026 Model H50 Overview

This section describes the features that come with the base configuration of a Model H50 and the optional features that can be used to enhance the Model H50's abilities.

Figure 22 shows the front view of a Model H50 in a S00 rack.



Figure 22. Model 7026-H50 - Front View in S00 Rack

4.5.1 Model H50 Standard Configuration

This section describes the standard configuration of a Model H50.

Table 42 outlines the features that comes in a standard configuration for a Model H50.

Table 42. Model H50 Standard Configuration

Model H50 Standard Configuration	
Microprocessor	332 MHz PowerPC 604e with X5 cache
Level 1 (L1) cache	32 KB data / 32 KB instruction

Model H50 Standard Configuration	
Level 2 (L2) cache	256 KB ECC
RAM (minimum)	128 MB ECC synchronous DRAM
Memory bus width	128-bit
Ports	One parallel, three serial, one keyboard, one mouse, 10 Mbps Ethernet, SCSI.
Internal disk drive	9.1 GB SCSI-2 F/W
Disk / Media bays	Thirteen (one used)/three (two used)
Expansion slots	Nine (seven PCI, two PCI/ISA)
PCI buses	Two 32 and one 64-bit
Memory slots	Two
CD-ROM drive	32X (Max) SCSI-2
Service Processor	Yes
Diskette drive	1.44 MB 3.5" diskette drive
SCSI adapters	Two integrated SCSI-2 F/W adapters
AIX operating system version	Version 4.3 (unlimited license is standard)
System dimensions	13.8" H x 17.5" W x 32.3" D (350 mm x 443 mm x 844 mm); 120 lbs (55 kg) The configuration will affect these numbers
Warranty	24 x 7, on-site for one year (limited) at no additional cost

4.5.1.1 Model H50 Reliability, Availability, Serviceability Features

The Model H50 features the same reliability, availability, and serviceability (RAS) features as the F50, with the addition of redundant hot-swap cooling and optional redundant and hot-swap power supplies.

The power supplies may be in either AC or -48v configurations.

4.5.2 Model H50 System Expansion

Table 43 shows the possible maximum processor, storage, and memory configurations.

Table 43. Model H50 System Expansion

Model H50 System Expansion	
SMP configuration	To 2, 3, or 4 processors
RAM	Up to 3 GB
Maximum Internal Disk Storage	Up to 227.5 GB

4.5.2.1 Model H50 Additional Memory

Either 128 MB (4 x 32 MB) or 256 MB (2 x 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. A standard configuration can consist of either a 128 MB (4 x 32 MB 200-pin ECC SDRAM DIMMs or 256 MB (2 x 128 MB)) of Error Checking and Correcting (ECC) SDRAM DIMMs. The memory has 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.

Figure 23 shows the rear view of the Model H50 and all its ports and slots.




Figure 23. Model 7026-H50 - Rear View

4.5.2.2 Model H50 Bays

The H50 provides three media and 12 internal disk bays. Figure 24 shows the internal media and disk bays of the Model H50. Media bays A1 and A2 are full-sized bays that provide access for removable media. A1 can also be used for a disk drive. B1 is used for the diskette drive. B2 is a hidden bay that can be used for a SCSI disk when required.




Figure 24. Model 7026-H50 - Internal Media Bays and Disks

Two media bays are filled by the 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 extra 9.1 GB SCSI IPL disk if the optional SSA (# 6428) or SCSI (# 6425) configuration is selected.

4.5.3 Model H50 Optional Features

This section describes the internal optional features that can be added to the Model H50 base configuration at an additional cost.

The status of a feature is indicative of these qualifications:

- A** To indicate features that are available and orderable on the specified models.

W Indicates that this feature is withdrawn.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards and cables, are not included.

Table 44 provides a list of features for the Model H50.

Table 44. Model H50 Optional Features

Feature Code	Description	Status
Processors		
4311	1-Way 332 MHz 604e Processor	A
4320/4338	2-Way 332 MHz 604e Processor	A
Memory		
9083	Base 128 MB SDRAM DIMMs	A
4107	64 MB (2x32) SDRAM DIMMs	A
4110	256 MB (2x128) SDRAM DIMMs	A
4106	256 MB (2x128) SDRAM DIMMs (select)	A
4093	Memory board for expansion (level 3)	A
Disk Bays		
6519	SCSI 6-pack hot-swap back plane	A
6539	SSA 6-pack back plane	A
Internal Disk Drives		
2913/2919	9.1 GB 1" Ultra SCSI Hot-Swap	A
2908	9.1 GB 1" Ultra SCSI Hot-Swap	A
3002	9.1 GB 10K RPM 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
3071	4.5 GB SSA Hot-Swap	W
3078	9.1 GB 10K RPM SSA Hot-Swap	W
3070	9.1 GB 10K RPM SSA Hot-Swap	A
3074	9.1 GB 1" SSA Hot-Swap	W
6509	Mounting hardware for 3.5" in 5.25" bay	A

Feature Code	Description	Status
Internal Tape Drives		
6156	20/40 GB 8 mm (Black)	A
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
Fibre Channel		
6227	Fibre Channel Adapter	A
SCSI Adapters		
6206	Ultra SCSI	A
6207	Ultra SCSI Differential	A
2494	Ultra SCSI 3-Channel RAID	A
6205	Ultra2 SCSI Long	A
SSA Adapters		
6222	SSA Fast-Write Cache Option	A
6230	IBM Advanced SerialRAID	A
6231	128 MB Cache Option	A
6235	32 MByte Fast-Write Cache Option Card	A
HIPPI Adapters		
2732	Serial Hippi Long-Wave Adapter	A
2733	Serial Hippi Short-Wave Adapter	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A

Feature Code	Description	Status
6310	ARTIC960RxD Quad Digital	A
6311	ARTIC960RxF Adapter	A
System Adapter		
2751	S/390 ESCON Channel PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
4959	Token-ring adapter	A
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2987	Ethernet AUI / RJ-45	A
4951	4-Port Ethernet	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
X.25 Adapter		
2961	X.25 Adapter (ISA)	W
Miscellaneous		
8741	Mouse, 3 button, black	A

Feature Code	Description	Status
6041	Mouse, 3 button, white	A
3752	Service Publications	A

4.5.3.1 Model H50 Internal Disk and Media 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.

The design feature that provides the upgrade 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.

There are three media bays and one hidden disk bay. One media bay contains the 32X CD-ROM, and another the diskette drive. The third media bay can hold a tape drive or a hard disk. Hard disks located in the media bays are not hot-swappable. They can be mirrored, if desired, by using the RAID mirroring function contained within AIX. If the available bays are occupied by a hard disk, any other optional devices, such as a tape backup 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.

4.5.3.2 Model H50 SCSI Adapters

The Model H50 disk subsystem incorporates dual integrated SCSI-2 Fast/Wide controllers. The standard configuration controls the first 6-pack 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 SCSI adapter through the cable provided in feature # 2447.

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).

4.5.3.3 Model H50 Graphics Accelerator

The Model H50 provides support for the GXT130P entry-level (# 2830) 2D graphics accelerator. 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 accelerator.

4.5.3.4 Model H50 Redundant AC Power Supply

The Redundant AC Power Supply (# 6296) is a modular power supply 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. -48 Volt DC is available as base select and additional power supplies.

4.5.4 Model H50 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:

- There can be a maximum of two graphics accelerators installed in the system.
- Select the features (# 6425 or # 6428) that include an SCSI boot disk configured in the media bay if both 6-packs using SCSI RAID or SSA disks are ordered.

4.5.5 Model H50 Publications

Table 45 provides the publications for the Model H50.

Table 45. Publications for the Model H50

Order Number	Title
SA38-0546	<i>RS/6000 Enterprise Server Model H50 User's Guide</i>
SA38-0547	<i>RS/6000 Enterprise Server Model H50 Installation and Service Guide</i>
SA38-0509	<i>Diagnostic Information for Multiple Bus Systems</i>
SA38-0516	<i>Adapters, Devices and Cable Information</i>
SA23-2652	<i>System Unit Safety Information</i>

Order Number	Title
SA38-0538	<i>PCI Adapter Placement Reference Guide</i>
SA23-2690	<i>Customer Support Information</i>

4.6 RS/6000 7026 Model H70 Overview

This section describes the major attributes of the Model H70, including the minimum configuration and optional features.

Figure 25 provides a view of a Model H70 from the front.




Figure 25. Model 7026-H70 - Front View

4.6.1 Model H70 Minimum Configuration

This section describes the Model H70 minimum configuration. Table 46 provides the attributes of a Model H70 minimum configuration.

Table 46. Model H70 Minimum Configuration

Model H70 Minimum Configuration	
Microprocessor	340 MHz RS64 II
Level 1 (L1) cache	64 KB data / 64 KB instruction
Level 2 (L2) cache	4 MB ECC per processor
RAM (minimum)	128 MB SDRAM
Memory bus width	128-bit
Ports	One parallel, three serial, one keyboard and one mouse, 10/100 integrated Ethernet, Power Controller Interface (PCI), SCSI port
Internal disk drive	9.1 GB Ultra SCSI
Disk/media bays	Thirteen (one used)/three (two used)
I/O expansion slots	Eight PCI
PCI bus widths	Four 32-bit and four 64-bit
Memory slots	Two
Power	220 volt AC; -48 volt DC; optional second power supply
Service processor	Yes
CD-ROM drive	32X (max) CD-ROM drive
Diskette drive	1.44 MB 3.5" diskette drive
Integrated SCSI controllers	Dual integrated Ultra SCSI controllers
AIX operating system version	Version 4.3 (unlimited user license is standard)
System dimensions and weight	13.8" H x 17.5" W x 34.3" D (350 mm x 443 mm x 870 mm); 175 lbs (71.4 kg) in a minimum configuration
Warranty	24 x 7, on-site for one year (limited) at no additional cost

Figure 26 shows the slot locations as seen from the rear view.




Figure 26. Model 7026-H70 - Rear View

4.6.2 Model H70 System Expansion

Table 47 shows the possible maximum processor, memory, and storage configurations.

Table 47. Model H70 System Expansion

Model H70 System Expansion	
SMP configurations	Two 2-way CPU cards
RAM	8 GB SDRAM
Maximum Internal disk storage	254.8 GB (with boot bay and media bay)

4.6.3 Model H70 Optional Features

This section describes the internal optional features that can be added to the Model H70 configuration at an additional cost.

The status of a feature is indicative of these qualifications:

- A** To indicate features that are available and orderable on the specified models.
- S** Indicates a feature that is supported on the new model during a model conversion; these features will work on the new model, but additional quantities of these features cannot be ordered on the new model.
- R** Indicates that the feature is not supported on the new model and must be removed during the model conversion.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, monitors, are not included.

Table 48 outlines the potential options that can be included in the Model H70.

Table 48. Model H70 Optional Features

Feature Code	Description	Status
Processors		
4317	1-Way 340 MHz RS64 II Processor	A
4319	2-Way 340 MHz RS64 II Processor	A
Memory		
9083	Base 128 MB SDRAM DIMMs	S
4107	64 MB (2x32) SDRAM DIMMs	A
4110	256 MB (2x128) SDRAM DIMMs	A
4106	256 MB (2x128) SDRAM DIMMs (select)	S
4093	Memory board for expansion (level 3)	R
4098	Memory board for expansion (level 5)	A
4119	512 MB (2x256) SDRAM DIMMs	A
Disk Bays		
6519	SCSI 6-pack hot-swap back plane	R
6426	SCSI 6-pack RAID select backplane	S
6539	SSA 6-pack back plane	A
6427/6428	SSA 6-packs, select	S
6547	Ultra SCSI 6-pack hot-swap back plane	A

Feature Code	Description	Status
Internal Disk Drives		
2911/3019	9.1 GB Ultra SCSI Hot-Swap	S
2913	9.1 GB 1" Ultra SCSI Hot-Swap	A
2908	9.1 GB 1" Ultra SCSI Hot-Swap	A
3002	9.1 GB 10K RPM 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
2909	18.2 GB 1" Ultra SCSI Hot-Swap	A
3101	18.2 GB 1" Ultra SCSI Hot-Swap	S
3072	9.1 GB SSA Hot-Swap	S
3070	9.1 10K RPM GB 1" SSA Hot-Swap	A
3074	9.1 GB 1" SSA Hot-Swap	A
6509	Mounting hardware for 3.5" in 5.25" bay	A
Internal Tape Drives		
6142	4/8 GB 4 mm	S
6147	5/10 GB 8 mm	S
6156	20/40 GB 8 mm (Black)	A
6154	20/40 GB 8 mm (White)	S
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2619	20X Speed CD-ROM	S
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
2838	GXT120P	S
SCSI Adapters		
6206	Ultra SCSI	A
6207	Ultra SCSI Differential	A

Feature Code	Description	Status
6208	SCSI 2 Fast / Wide	S
6209	SCSI 2 Fast / Wide Differential	S
2493	SCSI 2 Fast / Wide RAID	R
2494	Ultra SCSI 3-Channel RAID	A
6205	Ultra2 SCSI	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	S
6222	SSA Fast-Write Cache Option	A
6230	IBM Advanced SerialRAID Plus	A
6235	32 MB Fast-Write Cache Option Card	A
6231	128 MB Cache Option	A
Fibre Channel Adapters		
6227	Fibre Channel Adapter	A
HIPPI Adapters		
2732	Serial Hippi Long-Wave Adapter	A
2733	Serial Hippi Short-Wave Adapter	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
6310	ARTIC960RxD Quad Digital	A
6311	ARTIC960RxF Adapter	A
6309	Digital Trunk Quad Adapter	S
System Adapter		
2751	S/390 ESCON Channel PCI	A

Feature Code	Description	Status
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
4959	Token-ring adapter	A
2920	Token ring	S
2979	Auto LANStreamer token ring	S
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	S
2987	Ethernet AUI / RJ-45	S
4951	4-Port Ethernet	A
WAN Adapters		
2701	4-Port Communications Controller (ISA)	R
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
X.25 Adapter		
2961	X.25 Adapter (ISA)	R
Miscellaneous		
8741	Mouse, 3 button, black	A
6041	Mouse, 3 button, white	A

4.6.3.1 Model H70 System Processors

The available processors are RS64 II processors running at 340 MHz. One- or two-way cards are available to allow one- to four-way SMP configurations.

4.6.3.2 Model H70 Memory

There are two memory slots for a maximum memory capacity of 8 GB (using 256 MB DIMMs). Each memory card can be populated with eight DIMM pairs of 64, 256, or 512 MB SDRAM DIMMs.

4.6.3.3 Model H70 Internal Bays

There are four media bays on the Model H70. A minimum configuration would include a diskette drive, and CD-ROM, and two empty media bays onto which one may contain either a disk or tape drive and the other a disk. Disks installed in the media bays are not hot-swappable.

- Bay B2 is reserved for the hidden IPL disk.
- Bay B1 is dedicated for the 1.44 Floppy Drive.
- Bay A2 is a half-height bay dedicated for the CD-ROM (# 2624).
- Bay A1 is a general purpose half-height bay that could be used for tape or disk.

The media bays are connected to the first integrated Ultra SCSI controller on the I/O planar with the 4-drop SCSI cable as part of the base configuration. This cable must be terminated either with the # 2448 SCSI terminator or by connecting it to the optional first SCSI 6-pack.

There are two 6-packs that can be populated with 9.1 GB SSA disks or 9.1 GB and 18.2 GB Ultra SCSI disks. The 6-packs can also be populated with 4.5 GB disks if they are part of an F50 or H50 MES upgrade to an H70.

To the right of the media bays is the first 6-pack - Bank C. The bays are numbered one through six from bottom to top and referred to as bay C1, bay C2, and so forth. There are two ways to connect this backplane to the system if the customer has order # 2448. They can attach the first 6-pack to either the first or second integrated Ultra SCSI controller.

The first integrated Ultra SCSI controller will default to F/W speed to allow attachment to both the internal media bays and a first SCSI six-pack. If this controller is connected only to the media bays, the speed may be set to Ultra using SMIT.

The second 6-pack is numbered the same way as the first. It can be connected to the second Integrated Ultra SCSI Controller, or if that is in use,

it can be connected to one of the following features: Either the # 6208, supported only on upgrades or the # 6206 SCSI Adapter using SCSI cable # 2447 as part of the high- performance option.

Figure 27 shows the internal disk and media bays of the Model H70.




Figure 27. Model 7026 H70 - Internal Disk and Media Bays

4.6.3.4 Model H70 - SSA Option

Using the Advanced ServerRAID Adapter (# 6230), it is possible to boot from SSA in non-RAID configurations. However, the configuration tools still require a SCSI boot disk to be ordered.

One or two SSA backplanes can be ordered (# 6539) with the SSA bulkhead cable (# 2453), which provides a backplane to bulkhead connection and then a bulkhead to Advanced SerialRAID Adapter external ports connection. An Advanced SerialRAID Adapter (# 6230) must be ordered when there are two SSA backplanes. Both are connected in one loop and driven by the Advanced SerialRAID Adapter.

The cable option (# 2453) that must be ordered provides the internal 6-pack connection but exits the covers through the bulkhead plate connector at the back of the machine and connects to the external ports of the Advanced

SerialRAID Adapter (# 6230). This option also allows connection of up to two SSA 6-packs (12 SSA drives), all under the covers of the Model H70.

SSA RAID1 is supported through the software function of AIX.

4.6.3.5 Model H70 - Redundant Power Supply

An additional AC (# 6290) or DC (# 6299) may be ordered to provide fault tolerant redundant power. The first power supply is required; the second power supply is optional.

A system with two power supplies cannot have both an AC and DC power supply.

4.6.3.6 Model H70 - Slot Assignments

There are eight physical I/O slots on four I/O Channels:

- Two 64-bit slots at either 50 or 33 MHz
- Two 32-bit slots at 33 MHz
- One 64-bit slot at 33 MHz paired with one 32-bit slot at 33 MHz
- One 64-bit slot at 33 MHz paired with one 32-bit slot at 33 MHz

4.6.3.7 Model 7014-S00 Rack

The Model 7014-S00 Rack is the same rack that the Model H50 is normally installed in.

Every stand-alone H70 requires the rear anti-tip plate shipped along with the system to be installed. The Model H70 CPU drawer should be installed in the Model S00 rack starting from EIA position 32 (top) and progressing toward EIA position 1 (bottom). A maximum of four H70s can be installed in one 7014-S00 rack. A customer can also order a Model H70 and have it installed in the 7015-R00, with the difference being that the R00 is a white version of the black S00, and the R00 does not have the front door provided on the S00.

See 6.2, “Model 7014-S00 Overview” on page 226, for further information on the Model 7014-S00.

4.6.4 Model H70 RAS

The H70 has many design strengths that provide outstanding RAS characteristics, including:

- ECC memory and cache, parity detection on system, I/O, and PCI buses
- Hot-swap redundant cooling, environmental monitoring

- Hot-swap power supply and disk drives
- Available redundant power supply
- Repeat-Guard CPU deconfiguration when CPU error is detected
- Service processor, auto dial-out, programmable reboot

4.7 High-Availability Cluster Server Solution HA-H70

The section outlines the hardware and software products that are included in the HA-H70 solution.

4.7.1 Solution HA-H70 Overview

The HA-H70 is a two-node, high-availability cluster consisting of two Model 7026-H70s mounted in a 7014 Model S00 Rack with AIX Version 4.3 system software and HACMP 4.3 high-availability cluster software. One HA-H70 configuration includes the Model 7133-D40 SSA Disk Storage Subsystem. Additional storage options can be configured with the HA-H70.

A sample HA-H70 configuration includes:

- Two 7026-H70 Enterprise Servers
- One 7014-S00 System Rack
- One 7133-D40 Serial Disk system
- AIX Version 4.3 license for each H70
- Two HACMP 4.3 Enhanced Scalability features, including either ES, ESCRM, or HAS
- Redundant AC Power Supply in each H70
- Two PCI SSA Adapters and SSA cabling per H70
- Serial-to-Serial Port Cable for drawer/drawer connection (# 2936)
- Two 10/100 Mbps Ethernet adapters or two token-ring adapters per H70

The base HA-H70 offers configuration flexibility. Since the HA-H70 is comprised of Model 7026-H70, all H70-supported features are supported by the HA-H70. All optional features that are available for the H70 can be installed in this configuration.

4.7.2 Solution HA-H70 Configuration Notes

When planning or implementing an HA-H70, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

- There can be a maximum of two graphics accelerators installed in each H70 system.
- A system console or cluster administration station for the HA-H70 system is required.

4.8 RS/6000 7026 Model H80 Overview

This section describes the features that come with the base configuration of a Model H80 and the optional features that can be used to enhance the Model H80.

Figure 28 shows the front view of a Model H80 in a rack (second drawer) with its I/O drawer (first drawer).



Figure 28. Model 7026 H80 with I/O drawer - Front View in a Rack

Note

The diskette drive is accessible by removing the cover of the I/O drawer.

The Model H80 contains two units: One Central Electronics Complex (CEC), with mainly processors and memory, and normally one I/O drawer with the I/O related components. It is possible to get an additional I/O drawer at special request (RPQ 8A1140).

4.8.1 Model H80 Minimum Configuration

This section describes the minimum configuration of a Model H80.

Table 49 outlines the features that come in the minimum configuration for a Model H80.

Table 49. Model H80 Minimum Configuration

Model H80 Minimum Configuration	
Microprocessor	450 MHz (1-, 2-, or 4-way) or 500 MHz (6-way) RS64 III
Level 1 (L1) cache	128 KB data/128 KB instruction
Level 2 (L2) cache	2 MB (1-way), 4 MB (2-, 4-, or 6-way) ECC
RAM (minimum)	256 MB ECC synchronous DRAM
Memory bus width	128-bit
Standard ports	One parallel, four serial, keyboard, mouse, Ethernet 10/100BaseT
Internal disk drive	None (optional)
Disk / Media bays	0 ¹ / 2 (one used)
Expansion slots (hot-plug)	14 (10 3.3V at 66 MHz, 4 5V at 33 MHz)
PCI bus width	10 64-bit, 4 32-bit
Memory slots	Two
CD-ROM drive	32X (Max) SCSI-2
Service processor	Yes
Diskette drive	1.44 MB 3.5" diskette drive
SCSI adapters	SCSI-2 F/W (internal) Ultra2 SCSI (external VHDCI)
AIX operating system version	Version 4.3.3 ² unlimited user license
System dimensions (Central Electronics complex (CEC) or I/O drawer)	8.58" (5 U) H x 17.5" W x 32.3" D (218 mm x 445 mm x 820 mm); 90 lbs (41 kg) Weight will vary based on installed options
Warranty	On-site for one year (limited) at no additional cost

Model H80 Minimum Configuration

- ¹ Two optional bays with Dual Bay Mounting Kit
² AIX 4.3.3 and the APARs IY09047 and IY09814

4.8.1.1 Model H80 Processor Subsystem

The Model H80 allows a 1-, 2-, 4-, or 6-way configuration. For 1-, 2- and 4-way configurations, the processor cards are equipped with 450 MHz RS64 III processors. The 6-way configuration comes with 500 MHz RS64 III. Both variants of RS64 III processors have 128 KB of both instruction and data Level 1 (L1) cache. Concerning the L2 cache, all configurations, except the 1-way, have 4 MB of L2 cache (4-way set associative). The 1-way machine is only equipped with 2 MB L2 cache (4-way set associative).

4.8.1.2 Model H80 Memory

The base configuration consists of a 256 MB (on the 1-way CPU-board or the memory board) of Error Checking and Correcting (ECC) SDRAM DIMMs in a 1-way configuration. If there is more than one processor in the system, the memory must be installed in quads on a memory board; so, the base configuration for these systems consists of 512 MB (4 x 128 MB). A maximum of 16 GB of memory is supported when both memory boards are used.

Note

A one-way system is limited to only one memory card, which can be only used when the one-way processor card contains no DIMMs.

4.8.1.3 Model H80 Reliability, Availability, Serviceability Features

The Model H80 features the same reliability, availability, and serviceability (RAS) features as the Model F80, including the same hot-plug capability of the PCI slots.

4.8.1.4 Model H80 Enhanced Service Processor

The service processor is located on the primary I/O drawer and provides similar function as the service processor in the Model F80.

4.8.1.5 Model H80 I/O Drawers

The Model H80 is shipped with one I/O drawer, which provides 14 hot-plug PCI slots and all the I/O ports, such as serial, parallel, Ethernet, SCSI, keyboard, and mouse. Only the plugs for connecting the I/O drawer are in the Central Electronics Complex (CEC). The two units, CEC and I/O drawer, are connected with three kinds of cables: Remote I/O (RIO) cables (two per

system, three in case of two I/O drawers), JTAG cable (one per system), and the V/S COMM cable (CEC SPCN one per system).

As mentioned previously, it is possible to connect a secondary I/O drawer if more hot-plug PCI slots are necessary. The secondary I/O drawer can be added to the Model H80 system through RPQ 8A1140. It is not possible during the normal ordering process because of a performance impact in the RIO loop. The secondary I/O drawer provides 14 additional hot-plug PCI slots and two additional media bays. If more than one I/O drawer is used, then two additional System Power Control Network (SPCN) cables are necessary to connect the two I/O drawers. The I/O drawers used in the Model H80 are the same as in Model M80.

Figure 29 shows the view into the I/O Drawer.



Figure 29. I/O Drawer - Inner View

As shown in Figure 29, there are plastic separators between the hot-plug PCI slots for electrical isolation between the hot-pluggable PCI cards. The function of the hot-plug PCI is not only provided by the I/O drawer itself. The PCI card must also support a hot-plug function. Some cards cannot be

removed when the system is running even if they are hot-plug enabled, for example the SCSI card with the operating system disks connected to it.

4.8.1.6 Racks

For the CEC and the I/O drawer(s), you need either a S00, T00, or a T42 rack. Refer to Chapter 6, "The RS/6000 Server and Solution Racks" on page 223 for further information about the racks. A maximum of three Model H80s can be installed in a T00 or T42 rack.

Figure 30 shows the rear view of the I/O Drawer.

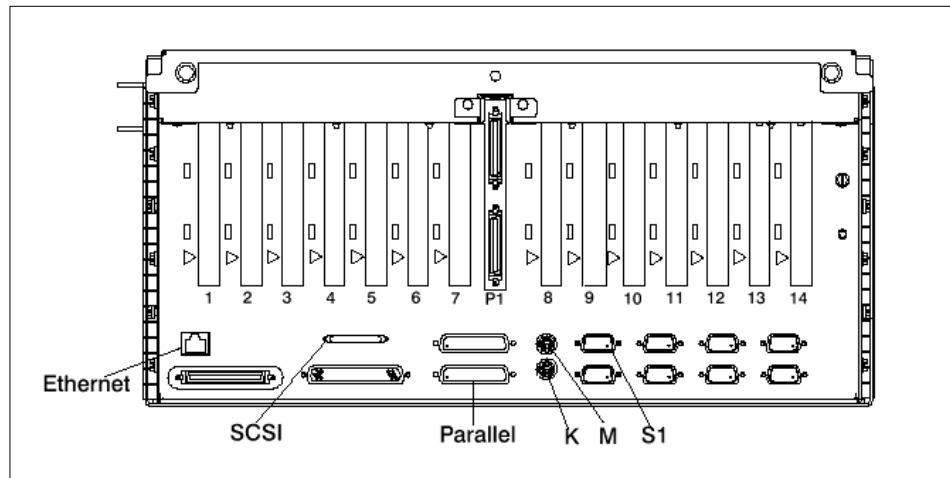


Figure 30. I/O Drawer - Rear View

Note

In Figure 30 there is an unlabeled port above the parallel port. This port is for diagnostics and is normally covert with a metal plate. It uses the same plug as the parallel port. To avoid confusion, this port should be always covert.

4.8.2 Model H80 Optional Features

This section describes the internal, optional features that can be added to the Model H80 base configurations at an additional cost.

The status of a feature is indicative of these qualifications:

- A** To indicate features that are available and orderable on the specified models.
- S** Indicates a supported device in this model.
- R** Indicates this adapter is not supported and must be removed during a model conversion from an F50.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards and cables, are not included.

Table 50 provides a list of the optional features available on the Model H80.

Table 50. Model H80 Optional Features

Feature Code	Description	Status
Processors		
5201	1-way 450 MHz RS64 III Processor	A
5202	2-way 450 MHz RS64 III Processor	A
5204	4-way 450 MHz RS64 III Processor	A
5206	6-way 500 MHz RS64 III Processor	A
Memory		
4075	Memory Board	A
4098	Memory Board, 16 Position	R
4107	64 MB SDRAM DIMM	A
9083	128 MB SDRAM DIMM	S
4110	256 MB SDRAM DIMM	A
4106	256 MB SDRAM DIMM	S
4119	512 MB SDRAM DIMM	A
4131	1024 MB SDRAM DIMM	A
Graphics Accelerators		
2838	GXT120P	S
2830	GXT130P	A
Power Supplies		
6282	Redundant AC Power Supply CEC	A

Feature Code	Description	Status
6283	Redundant AC Power Supply I/O drawer	A
6290	AC Power Supply, 750W	R
6299	DC Power Supply, 750W	R
Internal Disk Bays		
6540	SCSI Dual IPL Bay	A
6426	SCSI 6-pack	R
6425	SCSI 6-pack	R
6547	SCSI 6-pack Hot Swap	R
6427	SSA 6-pack	R
6428	SSA 6-pack	R
Internal Disk Drives		
2900	4.5 GB Ultra SCSI	S
2901	4.5 GB Ultra SCSI Hot-Swap	R
3071	4.5 GB SSA Hot-Swap	R
2908	9.1 GB Ultra SCSI	A
2911/3019	9.1 GB Ultra SCSI Hot-Swap	R
2913/2919	9.1 GB Ultra SCSI Hot-Swap	R
3002	9.1 GB Ultra SCSI 10K RPM	R
3008	9.1 GB Ultra SCSI Hot-Swap 10K RPM	R
3072	9.1 GB SSA Hot-Swap	R
3074	9.1 GB SSA Hot-Swap	R
3070	9.1 GB SSA Hot-Swap 10K RPM	R
3078	9.1 GB SSA Hot-Swap 10K RPM	R
2909	18.2 GB Ultra SCSI	A
3101	18.2 GB Ultra SCSI Hot Swap	R
3104	18.2 GB Ultra SCSI Hot Swap	R
Internal Tape Drives		

Feature Code	Description	Status
6142	4/8 GB 4 mm	S
6159	12/24 GB 4 mm	A
6147	5/10 GB 8 mm	S
6156	20/40 GB 8 mm (Black)	A
Internal CD-ROMs		
2619	20X Speed CD-ROM	S
2624	32X/40X Speed CD-ROM	A
SCSI Adapters		
6204	Universal Ultra SCSI Differential	A
6205	Ultra2 Dual Channel SCSI	A
6206	Ultra SCSI	S
6207	Ultra SCSI Differential	S
6208	SCSI 2 Fast / Wide	R
6209	SCSI 2 Fast / Wide Differential	R
2494	Ultra SCSI 3-Channel RAID	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	S
6222	SSA Fast-Write Cache Option	A
6235	32 MB Fast-Write Cache Option Card	A
6225	IBM Advanced SerialRAID	S
6230	IBM Advanced SerialRAID Plus	A
6231	128 MB DRAM Option Card	A
Fibre Channel Adapters		
6227	Fibre Channel Adapter	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
2948	ARTIC960Hx 4-Port T1/E1 PCI	R

Feature Code	Description	Status
2949	ARTIC960Hx DSP Resource	R
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
6309	Digital Trunk Quad Adapter	R
6311	ARTIC960RXF Digital Trunk Adapter	A
HIPPI Adapters		
2732	Serial Hippi Long-Wave Adapter	R
2733	Serial Hippi Short-Wave Adapter	R
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
System Adapter		
2751	S/390 ESCON Channel PCI	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
2920	4/16 token-ring adapters	S
2979	Auto LANStreamer token ring	S
4959	4/16 token-ring adapters	A
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	R
2987	Ethernet AUI / RJ-45	R
4951	4-Port 10/100 Mbps	A
WAN Adapters		

Feature Code	Description	Status
2962	2-Port Multiprotocol PCI	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
Miscellaneous		
8741	Mouse, 3 button, black	A
6041	Mouse, 3 button, white	A
3753	Options Library	R

4.8.2.1 Model H80 Additional Processors

Customers with an installed Model H80 may increase the number of processors in their system, up to a maximum of a 6-way SMP, by using feature conversions. Such conversions are accomplished by replacing the system's existing processor card with a processor card containing a greater number of processors. The existing Model H80 processor card being replaced is returned to IBM.

4.8.2.2 Model H80 Additional Memory

The Model H80 has two possible memory cards with 16 DIMM slots each. The memory cards must be filled in quads with the same size. Either 512 MB (4 x 128 MB # 4110), 1024 MB (4 x 256 MB # 4119), or 2048 (4 x 512 MB # 4131) SDRAM DIMMs (200 pin, 10 ns) can be added. When upgrading from an H70, it is also possible to use 32 MB DIMMs (# 4107). Memory feature # 4107 is available only on 7026-H70 to 7026-H80 model conversions to complete the H80 requirements for memory to be installed in quads when installed on the 16-Position Memory boards feature.

For a 1-way configuration, it is also possible to use the eight DIMM slots on the 1-way processor card. These slots can be filled in pairs instead of quads. If DIMMs are mounted on the processor card, it is not possible to install memory boards.

4.8.2.3 Model H80 Internal Disk and Media Bays

For booting from disk, external disk storage is required for the Model H80, as data storage internal to the drawers is not provided; however, optional boot bays are available in the I/O drawer. Boot support is provided through using SCSI adapters, SSA adapters, or from the network using Ethernet or token ring adapters. Externally attached 2104-DL1 SCSI or 7133-D40 SSA disk subsystems can be used to contain system boot disks, if desired.

One or two SCSI IPL disks can be mounted in the I/O drawer utilizing the Internal IPL Disk Mounting Hardware (# 6540) feature. The Internal IPL Disk Mounting hardware is available for use in the I/O drawer only and eliminates the use of PCI slots 13 and 14.

If you have higher RAS requirements, it is recommended that you consider attaching hdisk1 to another SCSI controller. Additional cables are required to accomplish this attachment.

- Option #1: Attaching hdisk1 to the integrated/external Ultra2 SCSI controller.

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
 - #2424; .6m Male-Male Ultra2 cable
 - #2118; .3m Female-Male Ultra2 cable
- Option #2: Attaching hdisk1 to a PCI Ultra2 adapter (f/code 6205) in the same/nearby RIO drawer (i.e. less than .6m away)

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
 - #2424; 0.6m Male-Male Ultra2 cable
 - #2118; 0.3m Female-Male Ultra2 cable
- Option #3: Attaching hdisk1 to a PCI Ultra2 adapter (f/code 6205) in a RIO drawer that is more than .6m away

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
- #2425; 2.5m Male-Male Ultra2 cable
- #2118; 0.3m Female-Male Ultra2 cable

In the primary I/O drawer, there is only one media bay that can be used for an additional drive, such as a tape drive. If a secondary I/O drawer is connected to the system, then two additional media bays are provided for use.

4.8.2.4 Model H80 Graphics Accelerators

For 2D graphics applications, the Model H80 offers the POWER GXT130P Graphics Accelerator. The POWER GXT130P is a versatile, low-priced graphics accelerator. It is a half-length PCI card with a standard 15-pin DIN output connector and can support multisync monitors with at least a 64 KHz horizontal scan capability or DDC-2B-compliant multisync monitors with at least a 38 KHz horizontal scan capability. For Model H80, it is a cost-effective choice for server environments that require a graphical user console.

The old GXT120P adapter is also supported for use, but this adapter is withdrawn from marketing and, therefore, not orderable.

4.8.2.5 Model H80 Redundant AC and DC Power Supply

System power supplies are available in AC or -48 Volt DC and are designed to support up to 6-way SMP systems with 16 GB of memory. An optional modular AC power supply can be added to provide redundant power for AC systems (# 6282 for CEC, #6283 for I/O drawer). Redundant power is standard on systems with DC power. The power supplies are also hot-plugged power supplies (in the CEC as well as in the I/O drawers) to allow repair when the system is running. Optional, uninterruptible power supply (UPS) systems are also supported for this system.

4.8.3 Model H80 Configuration Notes

When planning or implementing an RS/6000 Model H80, give special consideration to the following limitations, which are also listed in the announcement letter for this product:

- There can be a maximum of only one graphics accelerator installed in the system.
- External disk storage is required for the Model H80 as data storage internal to the drawers is not provided; however, optional boot bays are available in the I/O drawer. The internal boot disk option requires the space of two PCI slots (slot 13 and 14), leaving 12 slots available.
- A maximum of two H80, with one I/O drawer each, fit in a S00 rack, or three H80, with one I/O drawer each, fit in a T00 or T42 rack.

4.8.4 Model H70 Upgrades

The 7026-H70 systems converted to Model H80 systems will require replacement of all system processors. The first Model H70 processor card can be replaced with one Model H80 2-way, or greater, processor card through a feature conversion. This processor conversion is only available at the time of initial model upgrade. A maximum of one H70 processor card can be converted to one H80 processor card for each system being converted. The existing Model H70 processor being replaced is returned to IBM.

The following points should be considered:

- Memory DIMMs in Model H80 system must be installed in units of four (quads). Only in a one-way configuration is it possible to install four pairs on the one-way processor card.
- Only two internal disks can move to Model H80 if the additional two disk bay is installed in the system.
- AIX 4.3.3 with APARs IY09047 and IY09814 is required. Old system disks may be upgraded before using in Model H80. Rack 7015-R00 is not supported for Model H80.

4.8.5 Model H80 Publications

Table 45 provides the publications for the Model H80.

Table 51. Publications for the Model H80

Order Number	Title
SA38-0575	<i>RS/6000 Enterprise Server Model H80 Installation Guide</i>
SA23-2652	<i>System Unit Safety Information</i>

Table 52 provides the publications that may be optionally ordered.

Table 52. Additional Publications for the Model H80

Order Number	Title
SA38-0565	<i>User's guide</i>
SA38-0566	<i>Service guide</i>

The publications mentioned are also available for download from the Internet. If you do not know where to get the publications, ask your sales representative.

4.9 High-Availability Cluster Server Solution HA-H80

This section describes the hardware and software that completes the HA-H80 solution.

4.9.1 Solution HA-H80 Introduction

The HA-H80 is a two-node, high-availability cluster consisting of two RS/6000 Model H80s mounted in a 7014 Model T00 rack with AIX Version 4.3.3 or later system software and HACMP 4.3 or later high-availability cluster software. One HA-H80 configuration includes the Model 7133-D40 SSA Disk Storage Subsystem. Additional storage options can be configured with the HA-H80.

The HA-H80 solution package must include a minimum of the following hardware and software that must be ordered together, as specified, to quality for the HA-H80 package price:

- One central electronics complex (CEC)
 - One primary I/O drawer and attachment cabling
 - One HA Solution Indicator (# 0700)
 - One HA Solution Processor Card - Choose from:
 - One HA Solution Processor Card, 4-Way (# 0505)
 - One HA Solution Processor Card, 6-Way (# 0508)
 - 512 MB memory
 - Two PCI SSA adapters
 - Two LAN adapters
 - One redundant CEC AC or DC power supply (# 6282 or # 6285)
 - One redundant I/O drawer AC or DC power supply (# 6283 or # 6286)
 - One Serial-to-Serial Port Cable for Drawer/Drawer (# 3124)
Only one serial cable (# 3124) is needed for the two H80 servers. Redundant dedicated heartbeat and messaging paths are required via LAN and ASYNCH ports on each server.
 - One software preinstall (# 5005)
Redundant SSA or SCSI IPL disks are required. AIX and HACMP software will be preinstalled on 7133 SSA disks unless optional SCSI disks are ordered.
 - One AIX Version 4.3.3, or later license
 - One HACMP Version 4.3, or later license

Note

Additional optional 7026-H80 features can be added.

- One 7014-T00 system rack, including:
 - Two rack content specify: 7026-H80 (# 0128)
 - Two rack content specify: I/O Drawer (# 0176)
 - One rack content specify: 7133-D40 (# 0156)
 - One power distribution unit, side mount (# 6171, # 6173, or # 6174)

Note

Additional optional 7014-T00 features can be added.

- One 7133-D40 serial disk system, including:
 - Four SSA disk drives
 - Six advanced SSA cables
 - One RDS Rochester Manufacturing Integration (# 0987)

Note

Additional optional 7133-D40 features can be added.

4.10 RS/6000 7026 Model M80 Overview

This section discusses the RS/6000 Enterprise Server Model M80, the system expandability, and the upgrade path from Model H80.

The RS/6000 Model M80, as shown in Figure 31 on page 137, is a mid-range member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM. Positioned between the new Model H80 and the Model S80, the Model M80 provides 64-bit scalability using the 64-bit RS64 III processors that operate at 500 MHz and incorporate 4 MB of L2 cache per processor. With its two processor positions, the M80 can be configured into 2-, 4-, 6-, or 8-way SMP configurations using 2-way and 4-way cards. Memory can be expanded up to 32 GB in 1 GB, 2 GB, or 4 GB increments. The M80 also incorporates an I/O subsystem supporting 32-bit and 64-bit standard PCI adapters. The maximum configuration provides 56 available hot-plug PCI slots.



Figure 31. RS/6000 Enterprise Server Model M80

4.10.1 Model M80 Product Positioning

The M80 is targeted at the key commercial processing segments of e-business, ERP, SCM, and Business Intelligence. In each segment, the Model M80 can meet critical requirements. For example:

- In ERP, the M80 is an excellent application server with its powerful processors, memory capacity, and optional data storage capability.
- In e-business, the M80 can serve as a fast, highly reliable business-to-business Web server, going to other systems for business data or hosting the data storage itself.

4.10.2 Model M80 Description

The following sections briefly discuss the Model M80 and its features.

4.10.2.1 Central Electronic Complex and the I/O Rack

The Model M80 consists of a rack-mounted Central Electronics Complex (CEC) drawer containing the processors and memory connected to another rack-mounted drawer containing the media, hot-plug I/O slots, and optional boot bays. The CEC and I/O drawers, as shown in Figure 31, implement redundant power and redundant cooling. The M80 rack-mounted drawer

configuration offers flexibility regarding the number of CEC and I/O drawers that can be mounted in the rack, thus providing more compute and I/O power per square foot of valuable floor space. The M80 CEC requires 8U (EIA units) of rack space, while each I/O drawer requires 5U of space. Thus, a minimum M80 configuration requires only 13U of rack space, while a maximum configuration fits into a 28U space. Depending on the number of I/O drawers attached, up to two M80 systems can be installed in the 7014-S00 or the new 7014-T00 racks with room remaining to install external data storage. The maximum configuration consists of one CEC drawer and four I/O drawers.

The M80 primary I/O drawer integrates the following functional components:

- System support (service) processor
- Fourteen - PCI slots
- Two - Media bays
- One - 1.44 MB diskette drive (mounted behind a removable cover)
- Four - Serial ports
- One - Parallel port
- One - 10/100 Ethernet port
- One - Keyboard port
- One - Mouse port
- One - SCSI-2 F/W port for internal drawer components, such as media devices and (optional) IPL disks
- One - Ultra2 SCSI port for external attachment use
- The first media bay in the primary I/O drawer is reserved for the required CD-ROM. The second media bay may be utilized for any additional supported internal media device.

Each secondary I/O drawer provides an additional 14 PCI slots and two media bays.

4.10.2.2 Storage

External disk storage is required for the M80, as data storage internal to the drawers is not provided; however, optional boot bays are available in the first I/O drawer. The internal boot disk option requires the space of two PCI slots, leaving 12 available.

4.10.2.3 Reliability, Availability, and Serviceability (RAS)

The Model M80 provides CPU de-allocation function that enables the system to detect failing processors and take them offline without re-booting the system. The service processor can record this action and notify the systems administrator or service personnel of the condition. Computing operation continues with the processor deactivated, therefore, allowing repair to be scheduled at a time convenient to your processing schedule.

The hot-plug function of the PCI I/O subsystem brings new levels of up-time to your system for the replacement or addition of I/O adapters. AIX assists the user by guiding them through the process, thus helping to ensure success. As an alternative, an HACMP packaged solution with dual M80s, called HA-M80, is offered to provide industry-leading high availability.

A partial summary of the RAS features are as follows:

- Error checking and correcting (ECC) system memory and cache
- I/O link failure recovery
- Environmental sensing
- N+1 power
- N+1 cooling units
- Service processor for integrated system monitoring
- Concurrent diagnostics
- Dynamic Processor Deallocation

4.10.3 7026 Model H80 To Model M80 Conversion

Installed 7026 Model H80 systems can be converted to 7026 Model M80 systems. This conversion requires replacement of the entire CEC, including processors. The H80 I/O drawers (including internal media devices and SCSI IPL disks) and most PCI adapter cards carry forward to the upgraded M80 system. This conversion preserves your existing system serial number.

The model conversion consists of the following items:

- 7026-M80 Central Electronics Complex. The H80 CEC being replaced is returned to IBM.
- Support processor firmware update to recognize the M80 CEC.
- 7026-M80 labels preserving the customers serial number.
- M80 publications.

The following items are required and must be ordered separately:

- One - RS64 III 2-Way or 4-Way 500 MHz processor card (# 5200 or # 5203)
- 7026-H80 to 7026-M80 Model Conversion Processor and Exchanges
- One - memory board - 32 position (# 4073)
- Additional memory (if required) to meet the M80 1 GB minimum system memory

The following items may be carried forward from the predecessor system:

- Memory DIMMs.
- I/O Drawer attachment cables.
- Primary and secondary I/O drawers.
- Internal media devices.
- Internal SCSI IPL hard disks.
- PCI adapters (see list of options among the sections to follow and check for support of existing adapters).
- Additional optional features may be added as desired.

Note

The 7026-H80 to M80 upgrade is available for systems with either AC or DC power. A DC system upgrade requires changes to the DC rack power distribution. This requirement can be satisfied by ordering a new 7014-T00 rack with DC power or by contacting the IBM Center for Customized Solutions (CCS) to have an existing rack modified.

4.10.4 RS/6000 Model M80 Standard Configuration

Table 53 provides the Model M80 Standard configuration, and Table 54 provides the Model M80 system expansion capabilities.

Table 53. Model M80 Standard Configuration

Model S80 Standard Configuration and Features	
Processor 8U Rack Drawer	
Microprocessor	2-way 500 MHz RS64 III SMP
Level 1 (L1) cache	128 KB data / 128 KB instruction
Level 2 (L2) cache	4 MB per processor
RAM (minimum)	1 GB
Memory bus width	128-bit
I/O 5U Rack Drawer	
I/O slots	14 hot-plug PCI slots
I/O bus width	10 64-bit and 4 32-bit
I/O bus speed	10@ 66 MHz (3.3 v) / 4 @ 33 MHz (5 v)
Storage options	Boot capability from externally attached DASD drawers, or optionally, two internal boot disks, which require two I/O slots
Standard features	
Integrated ports	Keyboard, mouse, four serial, one parallel
Integrated bays	Diskette drive, CD-ROM, one bay available
Integrated controller	Ultra SCSI (internal, supports SCSI-2 F/W speeds), Ultra-2 SCSI (external), 10/100 Mbps Ethernet
Operating system	4.3.3 unlimited user licence

Table 54. Model M80 System Expansion

Model M80 Maximum Configuration	
SMP configurations	4- 6- 8-way 500 MHz
RAM	Up to 32 GB

Model M80 Maximum Configuration	
I/O	Up to three additional 5U rack drawers (42 hot-plug PCI slots and 6 bays additional)
Storage	IBM 2104 Expandable Storage Plus (Ultra-2 SCSI), IBM 7133 Serial Disk System, IBM 2105 Enterprise Storage Server

4.10.5 Model M80 Optional Features

This section describes the internal features that can be added to a configuration at an additional cost.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.
- S** Indicates a feature that is supported on the new model during a model conversion. These features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed.
- R** Indicates this adapter is not supported and must be removed during a model conversion from an F50.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 55 provides M80 optional features and their status.

Table 55. Model M80 Optional Features

Feature Code	Description	Status
Processors		
5200	RS64 III, 2-Way SMP, 500 MHz, 4 MB L2 Cache	A
5203	RS64 III, 4-Way SMP, 500 MHz, 4 MB L2 Cache	A
Memory		
4107	64 MB (2 X 32 MB) DIMMs, 10 NS SDRAM	A
9083	Base 128 MB (4 x 32 MB), DIMMs, 10 NS SDRAM	S

Feature Code	Description	Status
4106	256 MB (2 X 128 MB) DIMMs, 10 NS SDRAM	S
4110	256 MB (2 X 128 MB) DIMMs, 10 NS SDRAM	A
4119	512 MB (2 X 256 MB) DIMMs, 10 NS SDRAM	A
4131	1024 MB (2 X 512 MB) DIMMs, 10 NS SDRAM	A
4133	1024 MB Memory (8 X 128 MB DIMMs)	A
4134	2048 MB Memory (8 X 256 MB DIMMs)	A
4135	4096 MB Memory (8 X 512 MB DIMMs)	A
4073	Memory Board, 32-position	A
4075	Memory Board, 16-position	R
Host Attachment		
2751	S/390 ESCON Channel PCI Adapter	A
Internal Disk Drives		
2900	4.5 GB Ultra-SCSI 16-bit Disk Drive	S
2908	9.1 GB Ultra-SCSI 16-bit 1-inch Disk Drive	A
2909	18.2 GB Ultra-SCSI 1" High (25 mm) Disk Drive	A
6540	Internal IPL Disk Mounting Kit	A
Internal Tape Drives		
6142	4 GB / 8 GB 4 mm Internal Tape Drive	S
6147	5 GB / 10 GB 8 mm Internal Tape Drive	S
6156	20 / 40 GB 16-bit 8 mm Internal Tape Drive	A
6159	12 / 24 GB 4 mm Internal Tape Drive	A
Internal CD-ROMs		
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2830	GXT130P	A
SCSI Adapters		
6205	PCI Dual Channel Ultra2 SCSI Adapter	A

Feature Code	Description	Status
6204	PCI Universal Differential Ultra SCSI Adapter	A
2494	PCI 3-Channel Ultra2 SCSI RAID Adapter	A
6209	PCI Differential F/W SCSI Adapter	R
6208	PCI F/W SCSI Adapter	R
6207	PCI Differential Ultra SCSI Adapter	R
6206	PCI Ultra SCSI Adapter	R
SSA Adapters		
6225	IBM Advanced Serial RAID	S
6230	IBM Advanced Serial RAID PLUS	A
6235	32 MB Fast-Write Cache Option Card	A
6231	128 MB DRAM Option Card	A
6215	IBM SSA RAID Adapter	R
6222	Fast Write Cache Option	R
Async Adapters		
2943	8-Port Asynchronous Adapter EIA-232/RS-422	A
2944	128-Port Asynchronous Controller, PCI bus	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital Trunk Adapter	A
6311	ARTIC960RxD Quad Digital Trunk Resource	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
4959	Token-ring adapter	A
2979	Token-ring adapter	R

Feature Code	Description	Status
2920	4/16 token-ring adapter	R
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit Ethernet - SX PCI Adapter	A
2985/2987	Ethernet 10Base2/T/10Base5T	R
4951	4-Port 10/100 Mbps	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
Fibre Channel Adapters		
6227	Gigabit Fibre Channel Adapter	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A
HIPPI Adapters		
2732	Serial Hippi Long-Wave Adapter	R
2733	Serial Hippi Short-Wave Adapter	R
Power Related		
6283	Redundant AC Power Supply, I/O Drawer, 595 W	A
6284	Redundant AC Power Supply, CEC, 1100 W	A
6196	Processor Power Regulator, 76 A	A
6197	Processor Power Regulator, 190 A	A
6282	Redundant AC Power Supply, CEC, 645 W	R
Rack Related		
6160	9-track 1/2" Tape Drawer and PCI SCSI Adapter	S

Feature Code	Description	Status
6324	I/O Drawer, Primary, 5 EIA	A
6325	I/O Drawer, Secondary, AC Power, 5 EIA	A
6327	I/O Drawer, Secondary, DC Power, 5 EIA	A

4.10.5.1 Model M80 Configuration Notes

Keep in mind the following when configuring a model M80:

Processors

- The M80 system contains two processor slots. The available 2-way and 4-way processor cards can be combined to create 2-way, 4-way, 6-way, or 8-way systems.
- Installed systems may increase the number of processors by replacing the existing processor card with a more powerful processor or by adding an additional processor card if a processor slot is available.

Memory

- The 7026-M80 must contain a minimum of 1 GB of system memory. Memory is expandable up to a maximum of 32 GB.
- Memory is installed in groups of eight DIMMs mounted on 32-position memory boards. A maximum of two 32-position memory boards may be installed in a M80 system.
- All DIMMs in each set of eight must be the same capacity.
- A minimum of eight memory DIMMs must be installed on each 32-position memory board.
- The M80 memory is accessed through two related, but distinct, ports. Balancing memory across two memory cards allows memory accesses in a coordinated parallel manner and can be utilized to provide optimal performance. Configurations containing a single memory board will function properly, but the memory will only be accessed through one port and will not make use of the full memory bus bandwidth.

Power

- Redundant power supplies are required on M80 system configurations. AC powered systems have the first power supply for the CEC, and each I/O drawer is incorporated in their basic design. A second power supply must be ordered separately for each of these components to provide redundancy for enhanced system availability.

- A processor power regulator must be ordered for each system. Processor Power Regulator, 76A (# 6196) is required for 2-way and 4-way systems. Processor Power Regulator, 190A (# 6197), is required for 6-way and 8-way systems.

DC Powered Systems

- DC powered M80 systems operate at -48V and include redundant power supplies in both the CEC and I/O drawers.
- DC Powered M80 systems are supported only in the 7014-T00 rack equipped with DC power distribution panel # 6117. (# 6117 is a feature of the 7014-T00 rack.)
- DC Powered M80 systems are limited to a total of two I/O drawers, thus providing a total of 28 hot-plug PCI slots and four media bays.
- The 7014-T00 rack power distribution panel (# 6117) provides attachment for up to two M80 CEC drawers, two M80 I/O drawers, and four 7133-D40 SSA drawers. Up to two M80 systems can be mounted in a 7014-T00 rack if they are configured with one I/O drawer each. M80 systems configured with two I/O drawers are limited to one system per rack.
- Rack power distribution panel configurations and capacities are unique to support model specific system requirements. DC powered M80 systems may not be combined in racks with other DC powered systems, such as the 7026-H50, H70, or H80.

Racks

- The M80 systems CEC and I/O drawers are packaged as rack-mounted devices and are supported in the RS/6000 7014-T00, T42, and S00 racks. An existing 7014 rack can be utilized for AC powered M80s if sufficient space is available.
A maximum of two Model M80s can be installed in a T00 or T42 rack.
- The M80 I/O drawers must be placed a minimum of 4U (EIA Units) below the top of the rack for service purposes.
- Supported devices, such as SCSI or SSA disk subsystems or tape subsystems, may share the rack with the M80 system.

I/O Drawer Attachment

- The M80 system allows up to four I/O drawers to be attached to the system CEC drawer. It is recommended the I/O drawers be located in the same rack as the CEC drawer for service purposes; however, they can be mounted in separate racks if desired.
- I/O drawers are connected to the CEC drawer with the following cables:

- RIO cables for data transfer
- System control and initialization cable
- Power control cables
- RIO cable connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. Model M80 systems with non-looped configurations can experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.
- Two RIO loops are available, each supporting up to two I/O drawers. A minimum of two RIO cables are required to attach the first I/O drawer on each RIO loop. A third drawer to drawer RIO cable is required to complete the loop when an additional I/O drawer is attached to the loop. RIO cables are available various lengths (# 3142 = 3 meters, # 3143 = 6 meters, and # 3144 = 15 meters) to attach I/O drawers within a single rack or across multiple racks if desired.
- One System Control and Initialization (JTAG) cable (# 5992 = 3 meters or # 5993 = 6 meters) is required to connect the primary I/O drawer to the CEC.
- One CEC-to-Primary I/O Drawer power control cable (# 6132 = 3 meters or # 6136 = 6 meters) is required to connect the primary I/O drawer to the CEC.
- Power control for the secondary I/O drawers is provided from the primary I/O drawer via one loop. The number of Power Control (SPCN) cables required is equal to one plus the number of secondary I/O drawers attached to the system. Thus, a minimum of two power control cables (# 6006 = 3 meters, # 6007 = 15 meters or # 6008 = 6 meters) are required for attachment of the first secondary drawer to the primary drawer. Each additional drawer requires one additional power control cable to complete the loop attachment.

Disks, Media, and Boot Devices

- Boot support is supported via SCSI adapters, SSA adapters, or from network via Ethernet or token ring adapters.
- Externally attached 2104-DL1 SCSI or 7133-D40 SSA disk subsystems can be used to contain system boot disks if desired.
- One or two SCSI IPL disks can be mounted in the primary I/O drawer utilizing the Internal IPL Disk Mounting Hardware (# 6540) feature. The Internal IPL Disk Mounting hardware is available for use in the primary I/O drawer only and eliminates the use of PCI slots #13 and #14.

If you have higher RAS requirements, it is recommended that you consider attaching hdisk1 to another SCSI controller. Additional cables are required to accomplish this attachment.

- Option #1: Attaching hdisk1 to the integrated/external Ultra2 SCSI controller.

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
 - #2424; .6m Male-Male Ultra2 cable
 - #2118; .3m Female-Male Ultra2 cable
- Option #2: Attaching hdisk1 to a PCI Ultra2 adapter (f/code 6205) in the same/nearby RIO drawer (i.e. less than .6m away)

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
 - #2424; 0.6m Male-Male Ultra2 cable
 - #2118; 0.3m Female-Male Ultra2 cable
- Option #3: Attaching hdisk1 to a PCI Ultra2 adapter (f/code 6205) in a RIO drawer that is more than .6m away

Required cable feature codes:

- #3139; hdisk1 to RIO-drawer bulkhead
 - #2425; 2.5m Male-Male Ultra2 cable
 - #2118; 0.3m Female-Male Ultra2 cable
- The External SCSI Adapter to Internal IPL Disk Bays Cable Assembly (# 3139) provides a connection from internally mounted IPL disks to a Mini-68-pin VHDCI connector on the rear bulkhead of the I/O drawer. A converter cable (# 2118) is utilized to attach the VHDCI Mini-68-pin connector to a 68-pin P-style SCSI attachment cable for connection to the external SCSI port.
 - A CD-ROM is required for each system. Additional internal media devices are limited to one in the primary I/O drawer plus two in each secondary I/O drawer installed on the system.

Hot-Plug Options

- It is not necessary to power down the system to install certain hot-plug options. The following options are hot-plug capable:
 - CEC power supplies (assumes redundant power supplies installed)

- I/O drawer power supplies (assumes redundant power supplies installed)
- Most PCI adapters
- Any PCI adapter supporting the systems boot device or systems console is not hot-plugged.
- The PCI graphics accelerator, POWER GXT130P Graphics Accelerator (# 2830), is not hot-plug capable.

HA-M80 Solution

- Redundant power supply for the CEC and I/O drawer is not required on DC systems.
- Only one serial cable (# 3124) is needed for the two M80 servers. Redundant dedicated heartbeat and messaging paths required via LAN and Async ports on each server.
- Redundant SSA or SCSI IPL disks are required. AIX and HACMP software will be preinstalled on 7133 SSA disks unless optional SCSI disks are ordered.
- Additional optional 7026-M80, 7014-T00, and 7133-D40 features can be added.

4.10.5.2 Model M80 Publications

Table 56 provides information on the publications shipped with the Model M80.

Table 56. Publications Shipped with the Model M80

Order Number	Title
SA38-0576	<i>RS/6000 Enterprise Server Model M80 Installation Guide</i>
SA23-2652	<i>System Unit Safety Information</i>

4.11 High-Availability Cluster Server Solution HA-M80

The HA-M80 solution package must include a minimum of the following hardware and software that must be ordered together as specified to qualify for the HA-M80 package price:

- Two - 7026-M80 Enterprise Servers, each incorporating:
 - One - Central Electronics Complex
 - One - Primary I/O Drawer and Attachment Cabling
 - One - HA Solution Indicator (# 0700)

- One - HA Solution Processor Card 4-Way (# 0509)
- 2 GB Memory
- Two - PCI SSA Adapters
- Two - LAN Adapters
- One - Redundant CEC AC Power Supply (# 6 284)
- One - Redundant I/O Drawer AC Power Supply (# 6283)
- One - Serial to Serial Port Cable for Drawer/Drawer (# 3124)
- One - Software Preinstall (# 5005)
- One - AIX Version 4.3.3, or later license
- One - HACMP Version 4.3, or later license
- One - 7014-T00 System Rack, including:
 - Two - Rack Content Specify: 7026-M80 (# 0129)
 - Two - Rack Content Specify: I/O Drawer (# 0176)
 - Two - Rack Content Specify: 7133-D40 (# 0156)
 - One - Power Distribution Unit, Side Mount (# 6171, # 6173, or # 6174)
- Two - 7133-D40 Serial Disk Systems, each including:
 - Four - SSA Disk Drives
 - Six - Advanced SSA Cables
 - One - RDS Rochester Manufacturing Integration (# 0987)

Chapter 5. High-End Enterprise Servers

This chapter takes a look at the following S-Series RS/6000s:

RS/6000 Model S70 Advanced:

- 7013-S7A - Enterprise Server
- 7015-S7A - Enterprise Server
- 7017-S7A - Enterprise Server

RS/6000 Model S80:

- 7017-S80 - Enterprise Server

High Availability Advanced Cluster Solutions:

- HA-S7A
- HA-S80 with one or two racks

A description of the following models is given to put the models in a historical context:

- 7013-S70 - Enterprise Server - Withdrawn
- 7015-S70 - Enterprise Server - Withdrawn
- 7017-S70 - Enterprise Server - Withdrawn

Following these introductory model descriptions is an in depth description of the currently available systems and their features, the Model S70 Advanced, and the 7017-S80.

5.1 Introduction to the Servers

Designed for a broad range of applications serving medium to large businesses, RS/6000 Enterprise Servers come in Symmetric Multiprocessor (SMP) models that are well suited for mission-critical commercial, large e-business, or enterprise resource planning environments. The Model S7A provides an easy transition path to advanced 64-bit computing. The Model S80 is the latest member of the popular S-family of 64-bit enterprise servers.

5.1.1 Model S70

The Model S70 was shipped October 31, 1997 and was the first of a new generation of 64-bit, 4 to 12-way Symmetric Multiprocessing (SMP) Enterprise Servers. The S70 server provides the capacity and scalability for

linking mission-critical applications to a corporate intranet for the exploitation of evolving electronic business opportunities.

Effective December 13, 1999, IBM withdrew the Model S70 from marketing.

5.1.2 Model S70 Advanced

The S70 Advanced was launched October 23, 1998, and it is an enhanced model of the S70. The S70 Advanced can be used as a stand-alone server and can also be attached to the RS/6000 SP as an SP-attached server. Using a specially designed adapter (separately available) that fits within the S70 Advanced I/O drawer, the system is capable of connecting directly into the SP Switch fabric. The system is then ideally suited to handle large database transactions while allowing the other SP nodes to act as application servers. The S70 Advanced is managed and controlled just as a regular SP node using the SP-unique Parallel Systems Support Programs (PSSP) system's management software.

The RS/6000 Enterprise Server S70 Advanced is shipped and delivered with all the internal adapters and devices already installed and configured. The most current AIX Version 4.3 software is included with every S70 Advanced and may be pre-installed if desired.

The Model S70 Advanced is packaged in two side-by-side units. The first unit is the Central Electronics Complex (CEC), and it contains the 64-bit 262 MHz RS64 II processors and system memory. The second unit is a standard 19-inch I/O rack containing 10 EIA I/O drawers. Up to three more I/O racks can be added to a system for a total of four.

Installed RS/6000 J30/40/50 and R30/40/50 models can be upgraded to the S70 Advanced to provide both 32- and 64-bit application support (most 32-bit applications run unaltered on the S70). This capability should allow you to upgrade hardware independently of application software.

The 7013-S7A, 7015-S7A, and 7017-S7A are the same type of machine. However, to preserve the existing serial number, a model upgrade from the J30/40/50 and R30/40/50 to S70 Advanced requires that the old machine type is kept.

To upgrade from an S70 to an S70 Advanced, you must perform the following steps:

The following need to be removed as an MES upgrade:

- One Service Processor

- One (# 9295) Base SCSI I/O Drawer, 7 EIA

The following are feature codes that also need to be added as an MES upgrade. The removals and the additions Miscellaneous Equipment Specification (MES) need to be linked:

- One Upgrade from Model 7017-S70 to Model 7017-S7A
- One (# 6321) Primary I/O Drawer Group
- One (# 8320) Upgrade SCSI I/O Drawer, 10EIA
- One (# 8322) Support Processor Group
- One (# 8447) 16-bit PCI SCSI Adapter to 6-pack
- One (# 8519) Upgrade SCSI 6-pack Hot Swap Bays

To upgrade the processor books from 125 MHz to 262 MHz, the following features should be added to the order:

- For the first and third 4-way processor books:
 - (# 5316) RS64 II Processor, 4-way SMP, 262 MHz (MES only) (Right Hand)
- For the 2nd 4-way processor book:
 - (# 5317) RS64 II Processor, 4-way SMP, 262 MHz (MES only) (Left Hand)

5.1.3 Model S80

The Model S80 was announced on September 13th, 1999 and is a member of a new generation of 64-bit 6-, 12-, 18-, or 24-way symmetric multiprocessing (SMP) enterprise servers. The Model S80 can be used as a stand-alone server, but it can also be attached to the RS/6000 SP as an SP-attached server. When configured as an SP-attached server, the S80 is managed and controlled just as a regular SP node using the SP-unique Parallel System Support Program (PSSP) systems management software.

The RS/6000 Model S80 has rock solid design and leadership application performance that makes it a premiere solution for mission-critical ERP/SCM, Data Warehouse/Data Mart, OLTP, and e-business applications. It is the latest member of the popular S-family of 64-bit enterprise servers. The Model S80 blends new copper technology and enhanced backplane speeds with a proven system architecture to provide reliability you can count on. Scalability is additionally enhanced by doubling the number of processors to a total of 24, the amount of memory to 64 GB, and by leveraging the new performance enhancements of AIX Version 4.3.3. With the Model S80, you can manage

the evolution of your business into 64-bit computing while still supporting your existing 32-bit applications. In addition, the I/O subsystem supports 32-bit and 64-bit standard PCI adapters. Combined with an outstanding array of industry software and services, the Model S80 provides the power, capacity, reliability, and expandability to help move you into the next generation of mission-critical commercial computing.

The Model S80 is packaged as a Central Electronics Complex (CEC) and an I/O rack. The Model S80 Central Electronics Complex entry configuration starts with a 6-way scalable SMP, system that utilizes the 64-bit, 450 MHz, RS64 III processor with 8 MB of Level 2 (L2) cache per processor. The 6-way SMP can be expanded to a 24-way SMP and the system memory can be expanded to 64 GB. A fully configured system would be a 24-way processor with:

- 64 GB of system memory
- 56 PCI adapter slots (three of which are reserved for the service processor and the two required Ultra SCSI adapters)
- 48 disk bays that can be hot plugged
- 8 media bays (one of which is reserved for the CD-ROM drive)

The I/O Rack contains the first I/O Drawer with:

- A service processor
- A high-performance disk drive
- 32X maximum speed CD-ROM
- 1.44 MB 3.5-inch diskette drive
- Two PCI SCSI adapters

Up to three additional Model S80 I/O Drawers can be added. Additional I/O Racks can also be ordered with the Model S80. Existing RS/6000 7015 Model R00 and 7014 Model S00 racks can also be added for additional storage and communication drawers. This helps to protect your existing investment in SSA, SCSI, or Fibre Channel DASD.

The RS/6000 Enterprise Server Model S80 is shipped and delivered with all the internal adapters and devices already installed and configured. AIX Version 4.3.3 software is included with every Model S80 and can be preinstalled if desired.

5.1.4 Model S80 Upgrade Paths

Continued investment protection is demonstrated by the easy and attractive upgrade from existing Model S70 Advanced systems. Only from a Model S7A can you upgrade to a Model S80.

5.2 RS/6000 Model S70 Advanced Product Description

The following sections outline the standard and optional features of the S70 Advanced shown in Figure 32.




Figure 32. CEC and I/O Rack of the RS/6000 S70 Advanced

5.2.1 Model S70 Advanced Overview

In this section, the features of the S70 Advanced are described. The S70 Advanced contains two units connected by the Remote I/O (RIO) and System Power Control Network (SPCN) cables. The unit on the left (shown in Figure 32, is the Central Electronics Complex (CEC), and the other unit is the I/O rack.

5.2.1.1 Model S70 Advanced Service Processor

The service processor provides for excellent RAS service features, such as the call home on a hard failure causing a *system down* situation. In addition, the service processor provides functions, such as power-on/off of the system, reading the service processor and POST error logs, reading vital product data (VPD), changing the bootlist, viewing boot sequence history, and changing service processor configuration parameters. All of these can be performed remotely. Customers also have the capability to enable console mirroring on the system console so that they can monitor all remote console activity.

5.2.1.2 Dynamic Processor Deallocation

Another high-availability feature is dynamic processor deallocation. This feature allows the system to completely isolate a processor that has been determined to be operating below an acceptable level of reliability. To accomplish this, the system monitors processor status in real-time to determine when the number of recoverable errors has exceeded a specified threshold. When the threshold is exceeded, the system sends a report to the operating environment that a processor is unreliable. The AIX operating environment then re-routes all interrupts and processes to other processors and disables the unreliable processor complex without interrupting end user services. Firmware level 991216 must be installed on the Model S80. Since it is a feature that was introduced to AIX in February 7, 2000, the AIX Section of this book will discuss this function further.

The location of the URL that microcode or firmware can be downloaded from is:

<http://www.austin.ibm.com/support/micro/download.html>

5.2.1.3 The Service Director

Service Director generally automatically report problems based on default settings. The customer may modify the default values to prevent service director from placing a service call during hardware upgrades, testing, or in the event that the failed component is not covered by an IBM Service Agreement (for example, a third-party disk subsystem). The customer can also configure the product to only alert personnel within the customer's IT department (using e-mail). This function can be configured instead of, or in combination with, the ability to automatically place a service call to the IBM Service Center.

5.2.1.4 The Processor Subsystem

The RS/6000 Model S70 Advanced server initial configuration features one 4-way 262 MHz Power PC RS64 II book that can accommodate up to two

additional 4-way processor books. The processors include an L1 Cache split into a 64 KB instruction and 64 KB data. The L2 cache is an 8 MB ECC L2 cache for 262 MHz processors. If the Model S70 Advanced was ordered using an upgrade from a Model S70, there could be some of the RS64 processors remaining.

The Model S70 Advanced uses the 262 MHz RS64 II processor, which is based on the 125 MHz RS64 processor used in the RS/6000 Model S70 Enterprise Server. Both are 64-bit, PowerPC-compatible, four-way superscalar implementations. They are optimized for commercial workloads.

The RS64 II processor has five pipeline execution units: Branch, load/store, fixed point, complex fixed point, and floating point. The complex fixed point unit provides support for multiplication and division math. There is a dispatch buffer that can hold up to 16 current instructions, a technique that helps reduce latency. It also has an eight-deep branch buffer. The processor can sustain a decode and execution rate of up to four instructions per cycle.

All processor arrays use fault detection and correction techniques, such as redundancy, error checking and correction (ECC), parity, and retry. Together, these tools promote high reliability, availability, and data integrity. This enables full fault detection and correction coverage within the CEC.

The RS64 II features include:

- 64 KB on-chip L1 instruction cache
- 64 KB on-chip L1 data cache with one-cycle load-to-use latency
- 8.4 GB/s L2 cache bandwidth
- 32 byte on-chip data buses
- 262 MHz operating frequency
- 4-way superscalar design
- Five stage deep pipeline

Figure 33 on page 160 contains a diagram of the processor card layout used on the Model S70 Advanced. There are four processors per CPU card that share a system bus that connects to the memory controller complex.

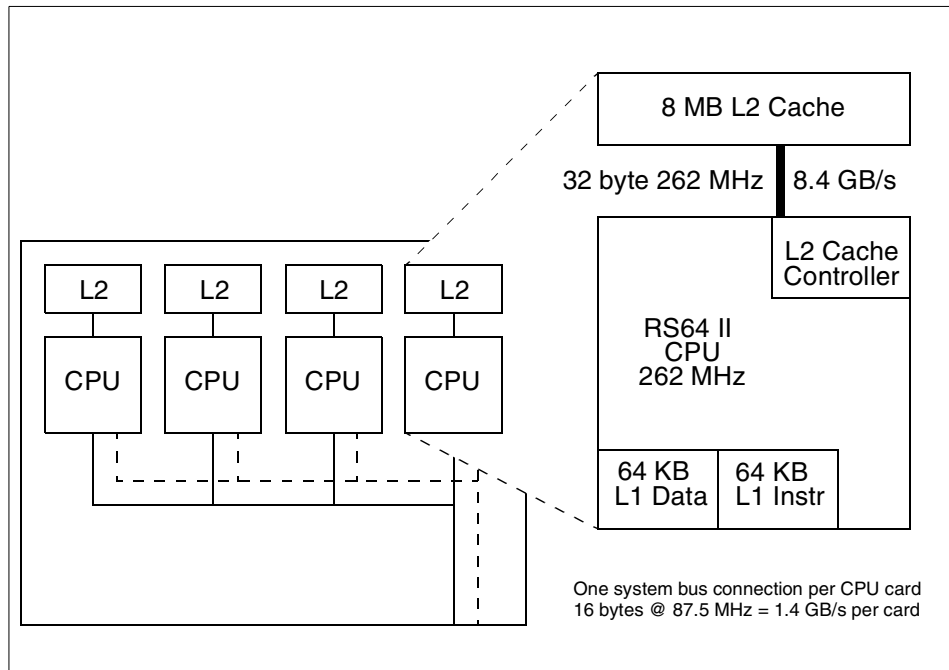


Figure 33. RS64 II Processor Card

5.2.1.5 Memory

The base configuration includes 1 GB of SDRAM-based memory. The order defaults to two # 4171 which are comprised of four 128 MB cards. The maximum current configuration is 32 GB. Memory sizes cannot be mixed within the four-card set. There is a maximum of 20 memory card slots.

The memory is accessed as two related, but distinct, ports. Balanced accesses use both ports in a coordinated parallel manner and can obtain up to twice the data in the same amount of time. Unbalanced configurations function properly, but the unbalanced portion of the memory is only accessed through one port and does not make use of the full memory bus bandwidth. For example, a system with 1 GB of memory should perform better if two 512 MB features are installed than if one 1 GB feature is installed.

Four system data buses link the Model S70 Advanced processor cards to the memory controller complex and the Remote I/O (RIO) hub. Each processor card has a single system data bus that connects to the data flow switch. The data flow switch connects to the memory subsystem through two memory ports. It also has a connection to the RIO hub. The switch consists of four data flow switch chips and a separate data flow control chip. The data flow

switch chips are the core of the memory controller. An SMP bus arbiter chip in the controller complex prevents switch chip contention.

Figure 34 contains a diagram of the layout of the memory controller complex.

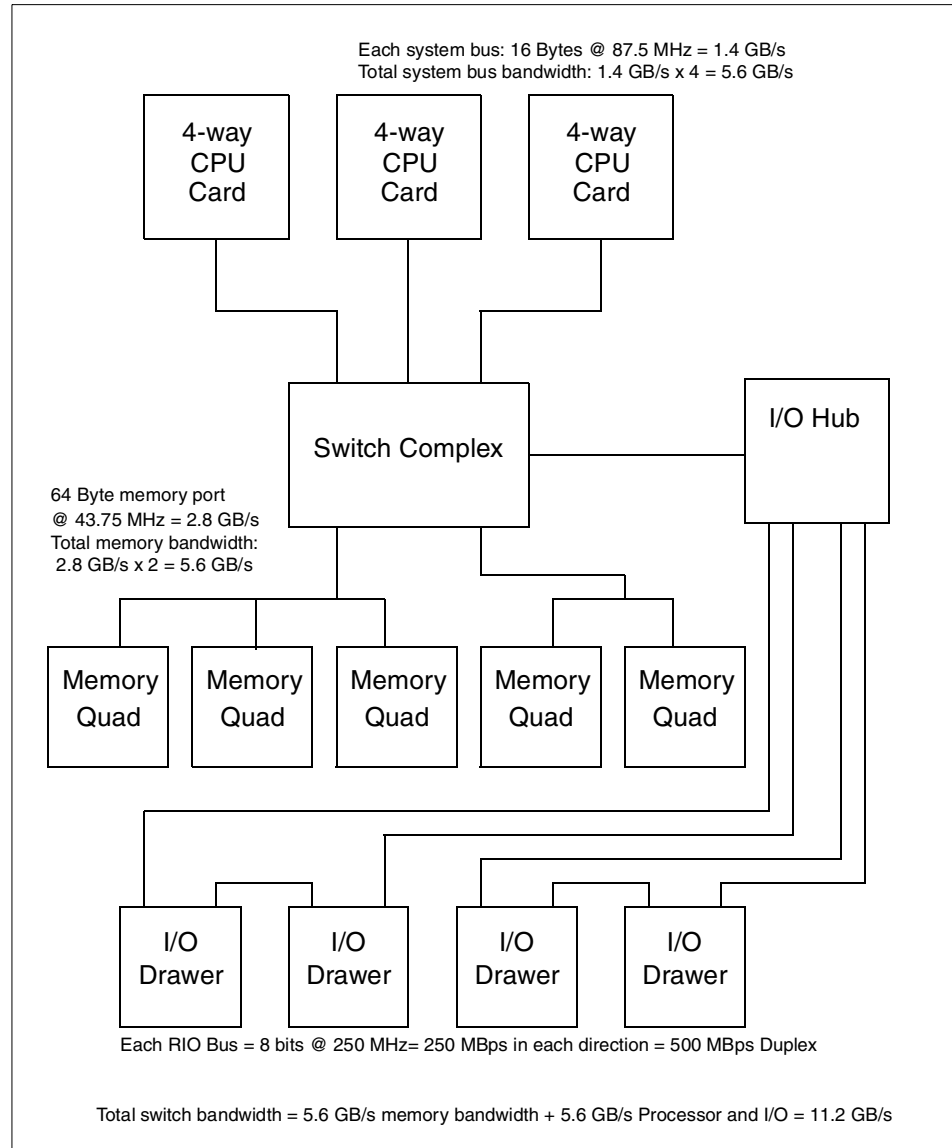


Figure 34. Model S70 Advanced System Switch Complex

The three system buses connecting the processor cards and the system bus connecting the RIO ports are 128 bits wide and run at 87.5 MHz to provide a total bandwidth of 5.6 GB/s. Memory ports are 512 bits wide and operate at 43.5 MHz. The two memory ports have an aggregate bandwidth of 5.6 GB/s. The total memory controller complex switch bandwidth is 11.2 GB/s.

The memory controller complex is mounted on a two-sided active backplane. The processors and memory are inserted as books. The I/O subsystem is connected to the complex using a set of RIO hub chips.

5.2.1.6 I/O Rack

There is a minimum of one I/O rack, and each I/O rack accommodates up to two I/O drawers (maximum four drawers available per system) with additional space for storage and communications subsystems.

I/O racks ordered as a feature of the S70 Advanced # 7000 must contain an I/O drawer. If additional external communication and storage devices (such as # 7133 and # 7027) do not fit in the space remaining in the S70 Advanced I/O racks, additional empty I/O racks should be ordered.

The base I/O drawer contains one high-performance 9.1 GB or two 4.5 GB Ultra SCSI disk drives, 32X (max) CD-ROM, 1.44 MB 3.5" diskette drive, two Ultra SCSI PCI Adapters, and a Service Processor, plus eleven available PCI Adapter slots, one media bay, and ten available hot-swappable disk drive bays.

The available rack is fitted with one I/O drawer, which contains the standard peripherals and SCSI controllers.

There is a maximum of 32 EIA units per I/O rack. Space for additional drawers in S70 Advanced I/O racks is limited to 22 EIA due to the I/O drawer (10 EIA).

5.2.1.7 I/O Drawers

There is a maximum of two I/O drawers (# 6320) per I/O rack. RIO cables and SPCN cables must be ordered for each additional drawer. Manufactured configuration of I/O drawers in I/O racks is based on cable lengths ordered. Refer to rack interconnection limitations for more information.

The primary I/O drawer must implement Ultra SCSI Adapters and SCSI 6-pack backplanes. Two or three Ultra SCSI SE Adapters are required depending on the desired configuration. The first SCSI Adapter controls only the systems media devices, such as CD-ROM and internal tape drives. Either one or two (# 6547) SCSI 6-Pack Hot Swap Backplanes may be installed in

the primary I/O drawer. Each backplane is controlled by a separate Ultra SCSI SE Adapter (# 6206) through an SCSI SE Adapter to a 6-pack cable (# 2447).

Figure 35 shows a 10 EIA I/O drawer when viewed from the front.




Figure 35. 10 EIA Drawer - Front View

Figure 36 on page 164 shows the I/O drawer rear view.




Figure 36. I/O Drawer Rear View (Partial) with 14 slots and 4 PCI Buses

Each drawer has four PCI buses per drawer: Slots 1-4 (PCI bus 1), 5-8 (PCI bus 0), 9-10 (PCI bus 3), and 11-14 (PCI bus 2). Slots 1, 5, 9, 10, and 14 are 64-bit slots. The remaining slots are 32-bit. The 32-bit adapters also function in the 64-bit slots. All slots are 33 MHz.

The standard peripherals required in the minimum configuration include the following:

- 1.44 MB Floppy.
- 32X Speed CD-ROM (# 2624)

The following feature code is the base I/O drawer:

- (# 6320) Base SCSI I/O Drawer, 10 EIA

The following lists the feature codes of the drawer groups:

- (# 6321) Primary I/O Drawer Group
- (# 6323) Secondary I/O Drawer Group

5.2.1.8 Power Distribution Units

The Model S70 Advanced is available with three Power Distribution Unit (PDU) options. They are selected with the following base feature codes:

- (# 9171) Side-Mounted, 1-Phase
- (# 9173) Side-Mounted, 3-Phase
- (# 9174) Side-Mounted, 3-Phase, Swiss

The following is a list of the feature codes of any additional Power Distribution Units:

- (# 6171) Side-Mounted, 1-Phase
- (# 6173) Side-Mounted, 3-Phase
- (# 6174) Side-Mounted, 3-Phase, Swiss

The S70 Advanced has six bulk power units (which are hot-swappable), in which five are in use, and the 6th is a stand-by.

The S70 Advanced has in the I/O drawer (# 6320) redundant power and cooling. From this I/O drawer, you will have two power cords that are both plugged either into the PDU or each cord can be plugged into its own PDU if the additional PDU is ordered, thus making this system highly available from a power point of view.

If you want this I/O redundant power option on an S70, the only way to get it is to upgrade the S70 to an S70 Advanced. You cannot get this I/O drawer and plug it into the S70.

5.2.1.9 System Power Control Network Cables

Additional power control cables (SPCN) must be ordered when extra I/O drawers or racks are ordered. The following feature codes are applicable to power control cables:

- (# 6006) Drawer-to-Drawer Control Cable
- (# 6007) Rack-to-Rack Power Control Cable

5.2.1.10 Remote I/O Cables

Similar to the power control cables, additional RIO cables are required when connecting extra I/O drawers or racks. RIO connections are always made in loops to help protect against a single point of failure resulting from an open, missing, or disconnected cable. S70 Advanced systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected during configuration, a problem is reported, but the system continues to boot.

Two RIO loops are available, each supporting two I/O drawers. If I/O drawers in the same I/O rack are to be connected in the loop, the 2-meter drawer-to-drawer cable (# 3126) must be ordered. The 6-meter base RIO cable will then be connected back to Port 1. If the I/O drawer in the rack is connected to the RIO Port 3 on the Central Electronics Complex, 15-meter rack-to-rack RIO cables (# 3127) should be ordered. When I/O drawers in the same RIO loop are in separate racks, 15-meter rack-to-rack RIO cables must be used. Manufacturing will determine the placement and cabling of I/O drawers based on the quantity of I/O racks and RIO cables.

Two 6-meter RIO cables are provided with the base system for connection to the base I/O drawer:

- (# 3126) Remote I/O Cable, Drawer-to-Drawer
- (# 3127) Remote I/O Cable, Rack-to-Rack

5.2.2 Model S70 Advanced Minimum Configuration

Table 57 provides the Model S70 Advanced minimum configuration.

Table 57. Model S70 Advanced Minimum Configuration

Model S70 Advanced Minimum Configuration	
Microprocessor	One 4-way 262 MHz PowerPC RS64 II CPU Card
Level 1 (L1) cache	64 KB data/64 KB instruction
Level 2 (L2) cache	8 MB per processor
RAM (minimum)	1 GB
Memory bus width	Dual 512-bit
Ports	One parallel, two serial, one keyboard, and one mouse
Internal disk drive	Two 4.5 GB (withdrawn) or one 9.1 GB Ultra SCSI (hot-swappable)
Media bays	Two (one available)
Expansion slots	Fourteen PCI (eleven available)
PCI bus width	32-and 64-bit
Memory slots	Twenty
CD-ROM Drive	32X (max)
Service Processor	Yes
Diskette Drive	1.44 MB 3.5-inch diskette drive
SCSI Adapters	Two Ultra-SCSI PCI Adapters
AIX operating system version	AIX 4.3.3 unlimited user license

5.2.3 Model S70 Advanced System Expansion

Table 58 provides the Model S70 Advanced system expansion capabilities.

Table 58. Model S70 Advanced System Expansion

Model S70 Advanced Maximum Configuration	
SMP configurations	Up to two additional 4-way processor books
RAM	Up to 32 GB
Internal PCI slots	Up to 56 per system (53 available)
Internal media bays	Up to eight per system
Internal disk bays	Up to 48 (hot-swappable)
Internal disk storage	Up to 436.8 GB

Note

There is no tape drive supplied in the standard configuration. Customers are able to select their preferred tape drive type as an additional feature.

5.2.4 Model S70 Advanced Optional Features

This section describes the internal features that can be added to a configuration at an additional cost.

The status of a feature is indicative of these qualifications:

- A** Indicates features that are available and orderable on the specified models.
- S** Indicates a feature that is supported on the new model during a model conversion. These features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed.
- W** Indicates features that are withdrawn from marketing on the specified models.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors, are not included.

Table 59 provides S70 Advanced optional features and their status.

Table 59. Model S70 Advanced Optional Features

Feature Code	Description	Status
Processors		
5317	Upgrade from 125 MHz - RS64 II 4-way 262 MHz exchange left	S
5316	Upgrade from 125 MHz - RS64 II 4-way 262 MHz exchange right	S
5313	RS64 II 4-way 262 MHz optional left	A
5312	RS64 II 4-way 262 MHz optional right	A
5314	Select RS64 II 4-way 262 MHz right	S
Memory		
4171	512 MB (4 x 128)	A
9168	Base 512 MB (4 x 128)	S
4173	1024 MB (4 x 256)	A
4174	1024 MB (4 x 256) select	S
4175	2048 MB (4 x 512)	A
4176	2048 MB (4 x 512) select	S
4177	4096 MB (4 x 1024)	A
4178	4096 MB (4 x 1024) select	S
4179	8192 MB (4 x 2048)	A
4180	8192 MB (4 x 2048) select	S
Host Attachment		
2751	ESCON Control Unit	A
8396	SP Attach Adapter	A
Internal Disk Drives		
2911/3019	9.1 GB Ultra SCSI Hot-Swap	S
2913	9.1 GB 1" Ultra SCSI Hot-Swap	A
2919	9.1 GB 1" Ultra SCSI Hot-Swap	S

Feature Code	Description	Status
3002	9.1 GB 1" 10K RPM Ultra SCSI Hot-Swap	A
3008	9.1 GB 10K RPM Ultra SCSI Hot-Swap	W
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
Internal Tape Drives		
6142	4/8 GB 4 mm	S
6147	5/10 GB 8 mm	S
6154	20/40GB 8 mm (White)	S
6156	20/40 GB 8 mm (Black)	A
6159	12/24 GB 4 mm	A
6160	9-Track 1/2" Tape Drawer	W
Internal CD-ROMs		
2619	20X Speed CD-ROM	S
2624	32X or 40X Speed CD-ROM	A
Graphics Accelerators		
2838	GXT120P	W
2830	GXT130P	A
SCSI Adapters		
6205	Ultra2 SCSI SE	A
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
6208	SCSI 2 Fast / Wide	S
6209	SCSI 2 Fast / Wide Differential	W
2493	SCSI 2 Fast / Wide RAID	S
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	W
6222	SSA Fast-Write Cache Option	A
6225	IBM Advanced Serial RAID	W

Feature Code	Description	Status
6230	IBM Advanced Serial RAID PLUS	A
6231	128 MByte DRAM Option Card	A
6235	IBM Advanced Serial RAID Cache Option	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital Trunk Adapter	A
6311	ARTIC960RxD Quad Digital Trunk Resource	A
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
4959	Token-ring adapter	A
2979	Auto LANStreamer token ring	S
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2986	Fast Etherlink XL 3Com	S
2987	Ethernet AUI / RJ-45	A
4951	4-Port Ethernet Adapter	A
WAN Adapters		
2962	2-Port Multiprotocol PCI X.25	A
Fibre Channel Adapters		

Feature Code	Description	Status
6227	Gigabit Fibre Channel Adapter	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A

5.2.5 Model S70 Advanced Configuration Notes

Keep in mind the following when configuring a model S70 Advanced:

- If you add a 7133-D40, an extra Power Distribution Unit (PDU) may be required.
- In an HACMP configuration, the standard native serial ports on the Model S70 Advanced are not available for heartbeat cabling. Therefore, you must install either the 8-port asynchronous adapter (# 2943) with serial-to-serial cable # 3125 or the 128-port Asynchronous controller (# 2944) with 128-port Asynchronous cable (# 8136) and rack-mountable Remote Asynchronous Node (# 8136) and the RJ-45 to DB-25 converter cable (# 8133) and the serial-to-serial cable (# 3125).
- Either an IBM supported ASCII terminal with an attachment cable or an IBM supported graphics display with an attachment cable and graphics accelerator is required for initial set up and must be available locally for service. OEM async. terminals are not recommended. They may not transmit the same character sequence and, thus, will not be recognized by the Service Processor.
- Model S70 Advanced SP attachment has the following notes:
 - The S70 Advanced can function as an attached SMP server within the IBM RS/6000 SP environment operating under the control of the IBM Parallel Systems Support Programs for AIX. Up to 16 systems may be attached in a single configuration. This interconnection can be accomplished utilizing the IBM RS/6000 SP System Attachment Adapter (# 8396) or through an Ethernet connection. Some I/O adapters available on the S70 or S70 Advanced systems are not supported in the SP environment and must be removed. Refer to the

IBM RS/6000 SP 9076-550 sales manual for a list of currently supported adapters.

- IBM Parallel System Support Programs for AIX Version 3.1 (5765-D51), or later, is required for an S70 Advanced system to function as an attached SMP server within the IBM RS/6000 SP environment.
- A minimum of one Ethernet adapter is required for an S70 Advanced system to function as an attached SMP server within the IBM RS/6000 SP environment. This adapter must be recognized by the S70 Advanced as en0.

Note

The RS/6000 SP System Attachment Adapter (# 8396) must always be located in slot # 10 of the primary I/O drawer. No adapters may be installed in slots # 9 or # 11 of the primary drawer when the SP attachment adapter is installed.

- If use of an EMC Fibre Channel Interface for AIX is required, the following must be done:
 - This RPQ (8A1135) is an I-Listed RPQ approval must be obtained before ordering. This RPQ delivers installation instructions (P/N 11K0808) for the EMC Fibre Channel Interface for AIX for use on the Model S70. The SRN for the adapter card is 910-XXX and shows a device name of *fchanx*. This information will NOT be found in *IBM Diagnostic Information for Multiple Bus*, SA38-0509. Refer to EMC Fibre Channel Interface for AIX kit for problem resolution. The FRU part number, 20K0445, can be ordered from IBM.
 - Ordering Instructions: Order the EMC Fibre Channel Interface for AIX directly from EMC. Contact your local EMC Sales Representative for EMC Kit CKIT-E70-AIX. This kit contains card P/N LP7000E-N1.

5.2.6 Model S70 Advanced Publications

Table 60 provides publications for the Model S70 Advanced:

Table 60. Publications for the Model S70 Advanced

Order Number	Title
SA38-0548	<i>RS/6000 Enterprise Server S70 Advanced Installation and Service Guide</i>
SA38-0549	<i>RS/6000 Enterprise Server S70 Advanced User's Guide</i>

Order Number	Title
SA38-0538	<i>PCI Adapter Placement Reference</i>
SA38-0509	<i>Diagnostic Information for Multiple Bus Systems</i>
SA38-0516	<i>Adapters, Devices, and Cable information for Multiple Bus Systems</i>

5.3 RS/6000 Model S80 Product Description

This chapter takes an in-depth look at the hardware packaging and features that comprise the RS/6000 Model S80. The standard and optional features of each model are described along with the features of the I/O Drawer and I/O Rack that are identical on both systems.



Figure 37. I/O Rack and CEC of the Model S80

5.3.1 Model S80 Processors

Each Model S80 processor card (or *book*) has six RS64 III processors with the associated L2 cache contained on the card. There are 8 MB of L2 per processor. Each processor card has the six processors on a set of two SMP system buses, and this dual bus interface is presented to the S80 backplane. All the processor cards in a system need to use the same type and speed of processor. The Model S80 can accommodate 4 CPU cards (4x6=24 processors).

The Model S80 uses the 450 MHz RS64 III processor, which is based on the 262 MHz RS64 II used in the Model S70 Advanced. Both are 64-bit, PowerPC-compatible, four-way superscalar implementations. They are

optimized for commercial workloads. The Model S80s processors, in particular, are designed to perform well running applications that place heavy demands on system memory. The RS64 III architecture addresses both the need for very large working sets and low latency. Latency is measured by the number of CPU cycles that elapse before requested data or instructions can be utilized by the processor.

The new processors combine IBM advanced copper chip technology with a redesign of critical timing paths on the chip to achieve greater throughput. The L1 instruction and data caches have been doubled to 128 KB each. New circuit design techniques were used to maintain the one cycle load-to-use latency for the L1 data cache.

L2 cache performance on the RS64 III processor has been significantly improved. Each processor has an on-chip L2 cache controller and an on-chip directory of L2 cache contents. The cache is four-way set associative. This means that directory information for all four sets is accessed in parallel. Greater associativity results in more cache hits and lower latency, which improves commercial performance.

Using a technique called Double Data Rate (DDR), the new 8 MB Static Random Access Memory (SRAM) used for L2 is capable of transferring data twice during each clock cycle. The L2 interface is 32 bytes wide and runs at 225 MHz (half processor speed), but, because of the use of DDR, it provides 14.4 GB/s of throughput.

Figure 38 on page 176 contains a diagram of the processor card layout used on the Model S80. There are six RS64 III processors per CPU card sharing two system buses that connect to the memory controller complex.

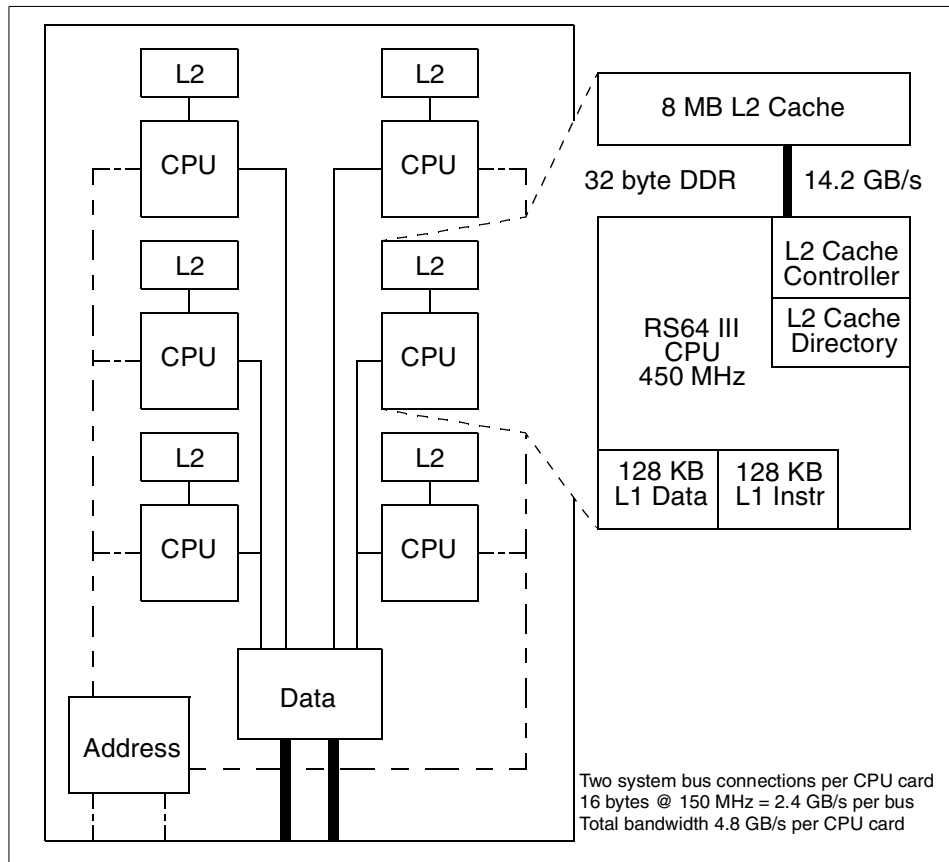


Figure 38. RS64 III Processor Card

The RS64 III processor has five pipeline execution units: Branch, load/store, fixed point, complex fixed point, and floating point. The complex fixed point unit provides support for multiplication and division math. There is a dispatch buffer that can hold up to 16 current instructions, a technique that helps reduce latency. It also has an eight-deep branch buffer. The processor can sustain a decode and execution rate of up to four instructions per cycle.

All processor arrays use fault detection and correction techniques, such as redundancy, error checking and correction (ECC), parity, and retry. Together, these tools promote high reliability, availability, and data integrity. This enables full fault detection and correction coverage within the CEC.

In summary, the RS64 III features include:

- 128 KB on-chip L1 instruction cache
- 128 KB on-chip L1 data cache with one cycle load-to-use latency
- On-chip L2 cache directory that supports up to 8 MB of off-chip L2 SRAM memory
- 14.4 GB/s L2 cache bandwidth
- 32 byte on-chip data buses
- 450 MHz operating frequency
- 4-way superscalar design
- Five stage deep pipeline

5.3.2 Model S80 Power Supply

The Model S80 comes equipped with sufficient power supplies for a 6- or 12-way system. When ordering a 12- or 24-way system, an additional power supply, processor regulator, and power regulator must be ordered:

- (# 6913) 1000 Watt AC Power Supply
- (# 6914) Programmable Power Regulator
- (# 6915) Processor Power Regulator

These additional power components are all installed in the front of the Model S80 CEC. The 1000 Watt AC Power Supply is installed in position P05; the Programmable Power Regulator is installed in position R08, and the Processor Power Regulator is installed in position M07. The locations are shown in Figure 39 on page 178.

5.3.3 Model S80 Memory

The Model S80 base configuration includes 2 GB of SDRAM-based memory. The maximum configuration is 64 GB. The Model S80 can accommodate up to 16 memory cards. Memory cards are used in sets of four, called quads. The memory subsystem provides ECC for single bit error correction and double bit error detection.

The Model S80 SDRAM memory cards contain redundant modules that support up to one in 14 memory modules not working. The memory is scrubbed by the controller, a feature that is designed to eliminate soft error problems. Memory cards are available in 512 MB, 1 GB, 2 GB, and 4 GB sizes. SDRAM modules are directly and permanently attached to the memory cards, a feature that minimizes failures and faults caused by connectors or sockets. A Model S80 must have a minimum of 2 GB of memory.

System memory is accessed through four related but distinct ports. A system should be configured with a minimum of two memory quads to make best use of the system architecture. Multiple ports can be accessed in a coordinated parallel manner and can obtain more data in the same amount of time. A configuration that uses only one port will function properly, but the system cannot make use of the full memory bus bandwidth. For example, a system with 2 GB of memory will perform better with two 1 GB features installed than if one 2 GB feature is installed.

The physical layout of the front of Model S80 CEC is shown in Figure 39.

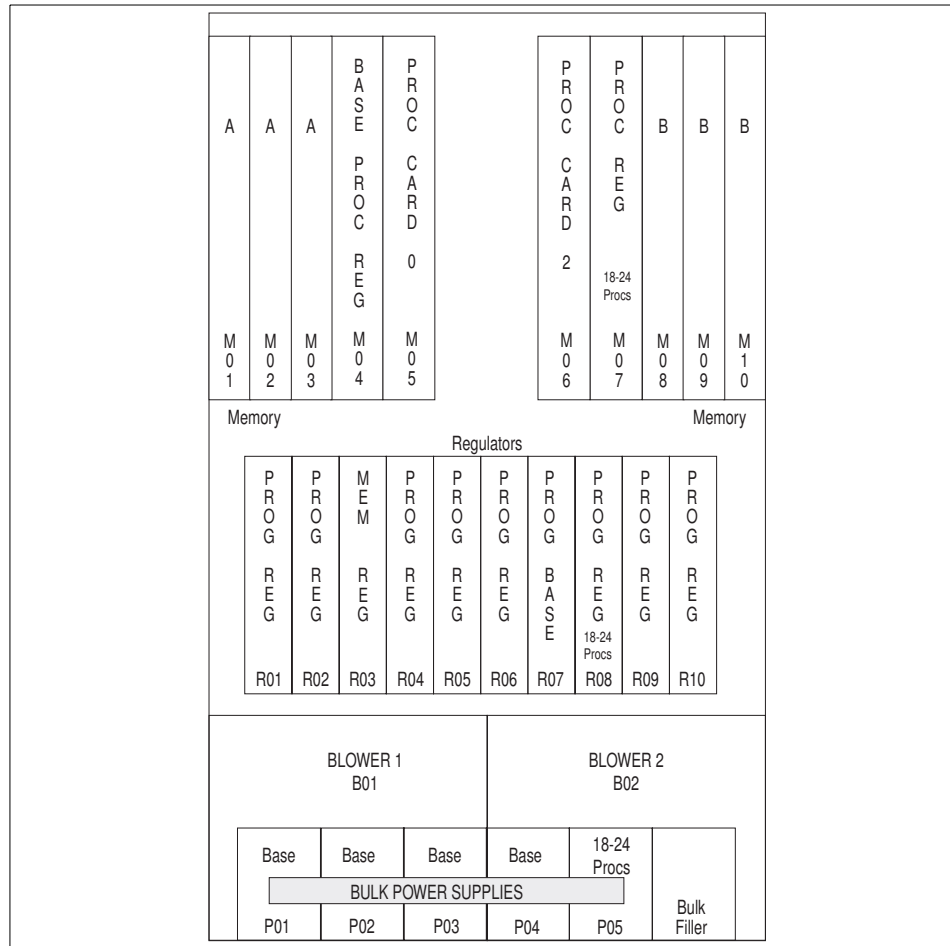


Figure 39. Front View of Model S80 CEC

The memory quads on the system are referred to as A, B, C, and D. Three memory cards, each from quads A and B, are installed on the front of the system. The remaining single card from each quad is installed in the rear of the system. Figure 40 shows the rear of the Model S80 CEC. All four cards of each of the C and D quads are installed in the rear of the system.

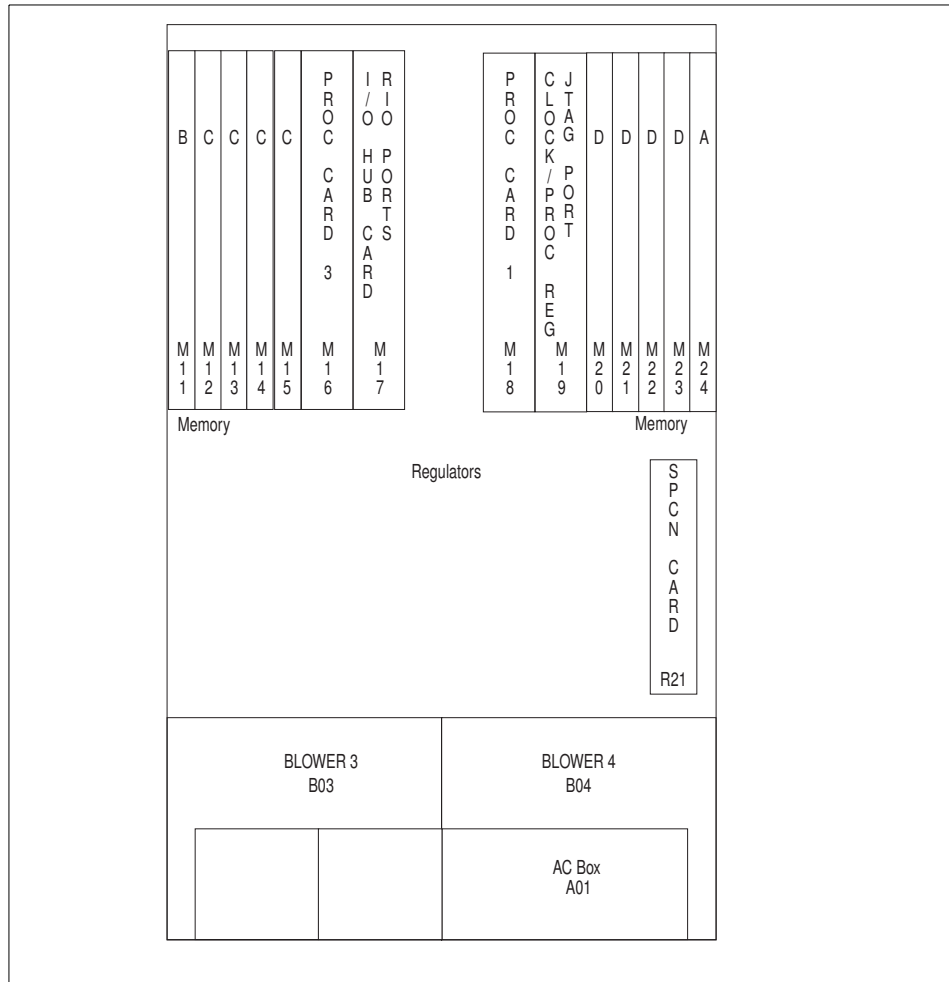


Figure 40. Rear View of Model S80 CEC

5.3.3.1 Model S80 Memory Controller Complex

Ten system data buses link the Model S80 processor cards to the memory controller complex and the RIO hub. Each processor card has two buses that connect to two interconnected data flow switches. Each switch connects

directly to two quads of Synchronous Dynamic Random Access Memory (SDRAM). Each switch also connects to its own dedicated port on the RIO hub.

Each switch consists of four data flow switch chips and a separate data flow control chip. The data flow switch chips are at the core of the memory controller. An SMP bus arbiter chip in the controller complex prevents switch contention.

Figure 41 on page 181 contains a diagram of the memory controller complex layout. The eight system buses connecting the processor cards and the two system buses connecting the RIO ports are 128 bits wide and run at 150 MHz to provide a total bandwidth of 24 GB/s. Memory ports are 512 bits wide and operate at 75 MHz. The four memory ports have an aggregate bandwidth of 19.2 GB/s. The total memory controller complex switch bandwidth is an impressive 43.2 GB/s. Transfer buffers are used in the switch to queue traffic if the needed connections are not immediately available.

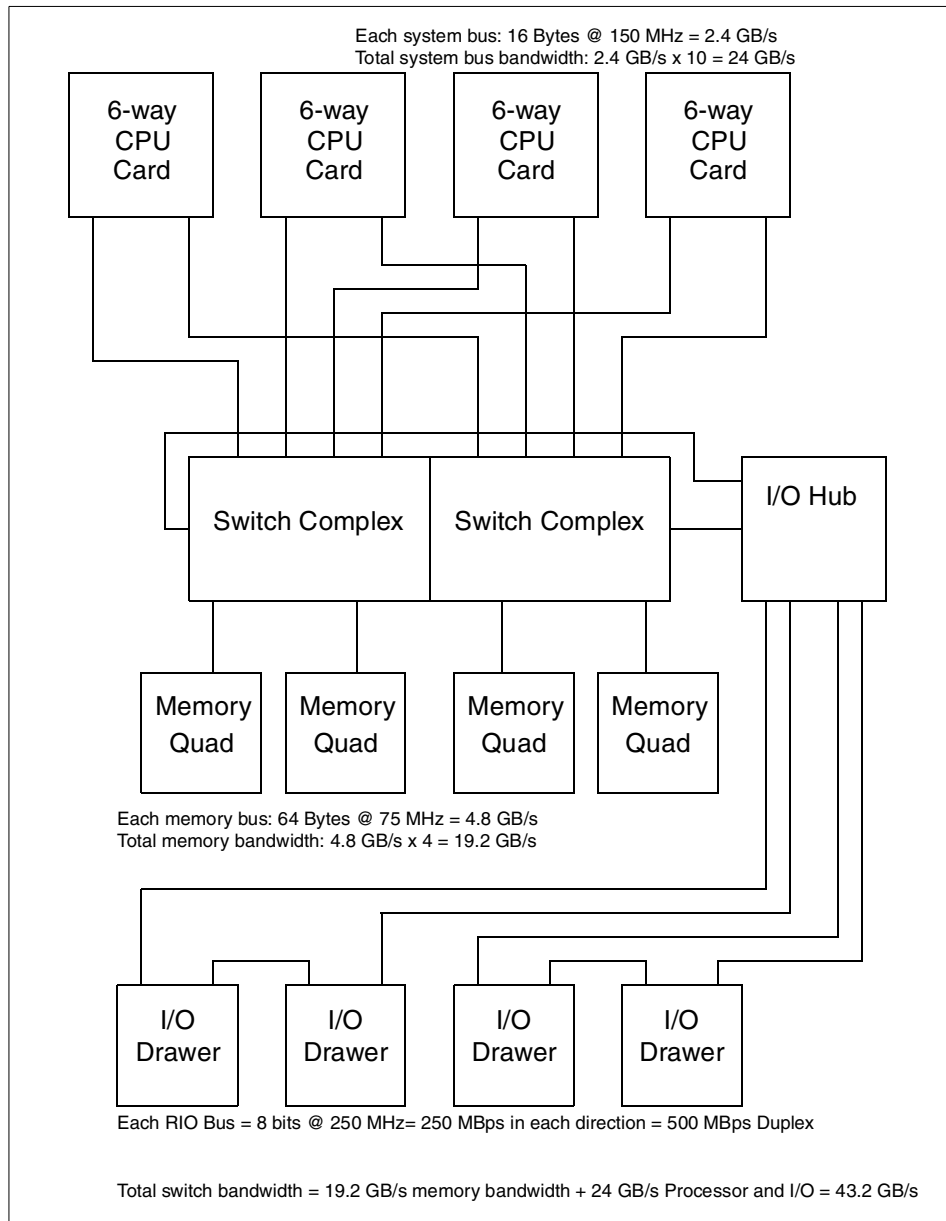


Figure 41. Model S80 System Switch Complex

In addition to the data flow switches, the design incorporates high-function address and data buffers to minimize latency. Addressing is done via a

separate 64 bit path. Cross-port traffic between processor cards is queued at the switch as needed.

The high-function address and data buffer chips break logical buses into smaller physical units. This allows the clocking frequency and, thus, the transfer speed to be increased. The buffers allow each of the system buses to support up to 2.4 GB/s of throughput.

The memory controller complex is mounted on a two-sided active backplane. The processors and memory are inserted as books. The I/O subsystem is connected to the complex via a pair of new RIO hub chips. These chips are on a replaceable I/O interface card, which will make upgrades easier in the future.

One of the features of this memory complex design is its ability to handle the next generation of processors. The system is designed to accommodate upgrades without losing its overall excellent balanced performance.

5.3.4 Model S80 Cabling

The CEC and the I/O Drawers are connected by various cables. The primary I/O Drawer has additional connections.

5.3.4.1 Model S80 System Power Control Network

The function of the SPCN allows a single switch on the front of the system CEC to control power to all of the I/O Drawers.

All I/O Drawers and the CEC must be connected in a single SPCN loop. The SPCN can function with any single connection broken, regardless of the location of the open connection.

The rules for cabling the SPCN are the same for all S-Series servers. A minimum of two SPCN cables is required for attachment of the first I/O Drawer on each system. Each additional I/O Drawer requires one additional SPCN cable for loop attachment. The available SPCN cable lengths depend on the Model of S-Series server. Figure 42 on page 183 shows the minimal SPCN loop configuration.




Figure 42. Base SPCN Loop with One I/O Drawer

Subsequent I/O Drawers are added to the SPCN loop. The length of SPCN cable chosen will depend on whether the drawer being added is in the same I/O Rack as an existing drawer. Figure 43 shows the maximum SPCN configuration.




Figure 43. Fully Configured SPCN Loop with Four I/O Drawers

The available SPCN cables for Model S80 systems are as follows:

- (# 6006) 2-meter drawer-to-drawer control cable
- (# 6007) 15-meter rack-to-rack control cable
- (# 6008) 6-meter rack-to-rack control cable

5.3.4.2 Model S80 Remote I/O Cables

On the Model S80, the Remote I/O (RIO) subsystem connects to the CEC via the I/O hub chips on a new replaceable I/O interface card. On the Model S70 Advanced, the chips are mounted directly on the backplane. In either case, four RIO connections are supported allowing a maximum of two RIO loops. The RIO connections are scalable, high-speed, point-to-point interfaces designed for low-latency high-bandwidth connections between two boards or boxes. Each RIO bus supports up to 500 MBps total or 250 MBps in each direction concurrently. The loop connection provides redundant paths; so, if a failure occurs in part of a cable, a warning message will be displayed, but the system will continue to operate.

RIO cables connect the CEC to the I/O devices located in the I/O Drawers. The RIO connections are set up as loops. The I/O hub chips direct the traffic around the loop in an optimal way for performance, and they redirect traffic if there are link errors.

The new RIO hub interface chip on the Model S80 offers improved buffering to enhance the effectiveness of the I/O interface.

These RIO connections are the key to allowing an expandable number of I/O Drawers that are physically separated from the CEC. In turn, this feature also enables the high number of PCI buses and slots.

The I/O Drawers are connected to the system CEC by a number of different cables. The Remote I/O (RIO) cables provide the means by which data can be transferred between the CPUs, and memory in the CEC and the storage and network devices can be connected via PCI adapters in the I/O Drawers.

The lengths of available RIO loop cable depend on the system being configured. The RIO cables used for Model S80 systems have different part numbers and feature codes than those used on Model S70 Advanced systems. RIO cables from Model S70 Advanced systems must not be used on Model S80 systems.

The primary I/O Drawer must be installed and connected to RIO port 0 of the CEC. The connection must be made from RIO port 0 of the CEC to RIO port 0 of the primary I/O Drawer. This connection is required to make the primary

drawer the first drawer in the loop that allows the firmware to initialize the system. It also allows the system to find the boot device.

An RIO loop supports up to two drawers. Each S-Series server supports up to two RIO loops. The configuration for a single I/O Drawer is very simple. There can only be one RIO loop since there is only one I/O Drawer. An example of this is shown in Figure 44.

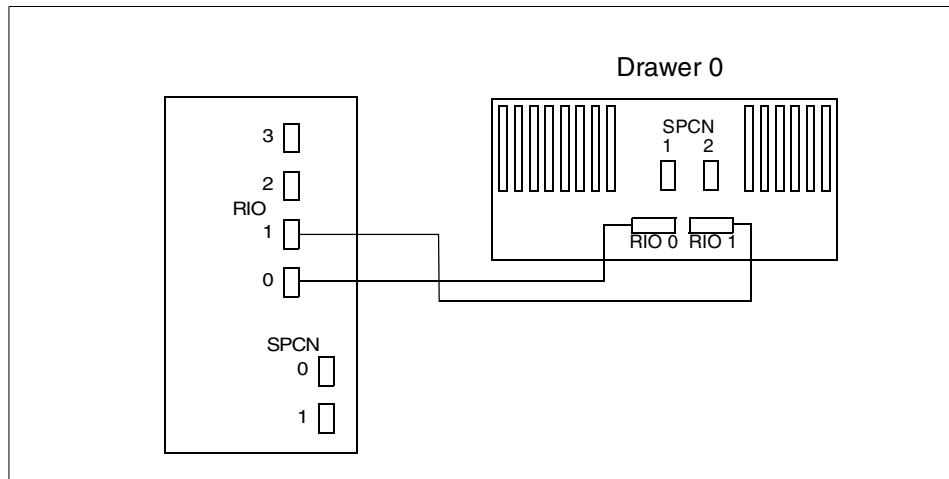


Figure 44. RIO Loop with One I/O Drawer

When adding a second drawer, you can cable the additional drawer into the existing RIO loop. This requires one additional RIO cable. An example of this is shown in Figure 45 on page 186.

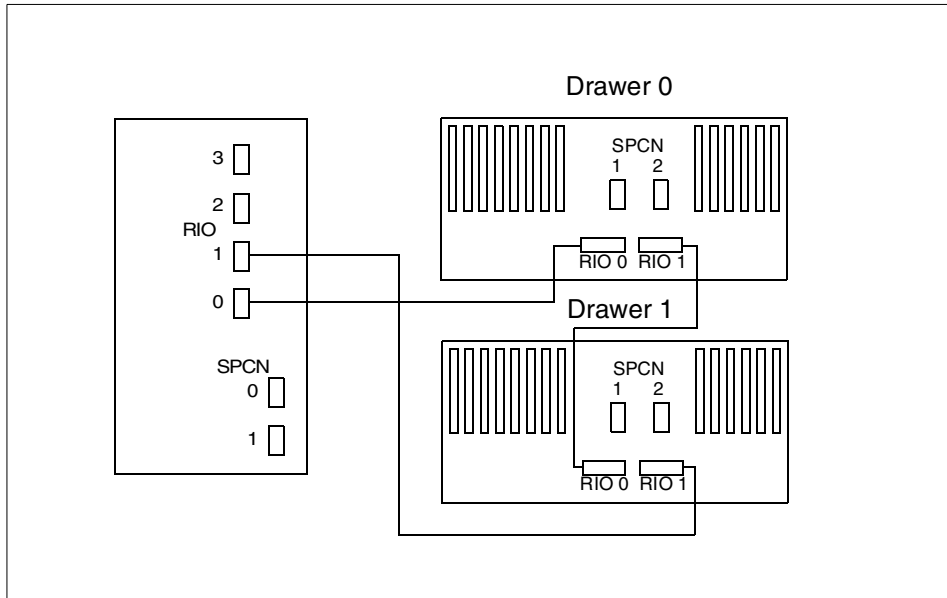


Figure 45. One RIO Loop Attached to Two I/O Drawers

Alternatively, you can cable the second I/O Drawer on its own loop. This option requires two additional RIO cables and provides each drawer with 500 MB of bandwidth. An example of this is shown in Figure 46 on page 187.

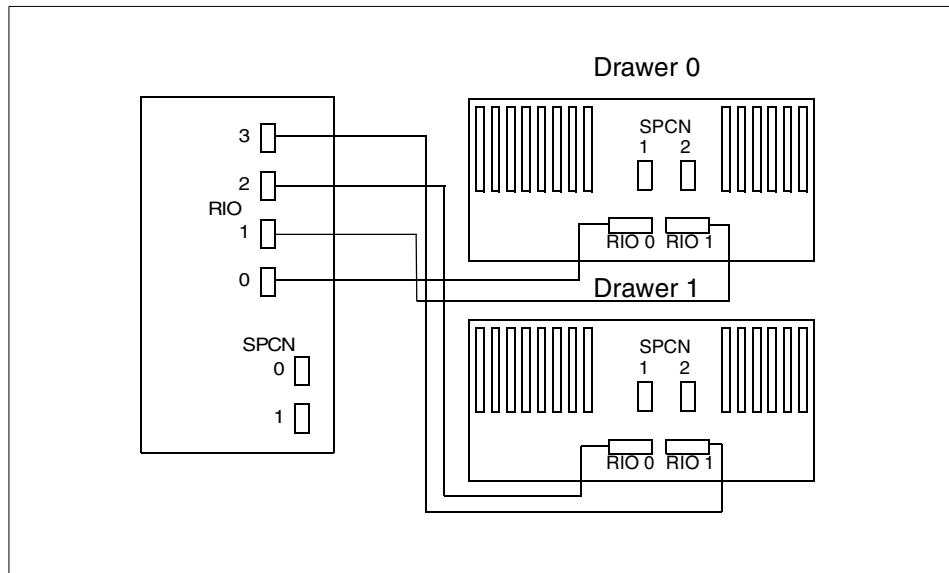


Figure 46. Two RIO Loops Attached to Two I/O Drawers

When adding a third I/O Drawer to a system, the number and length of RIO cables required will depend on the existing RIO loop configuration. If only one RIO loop is configured, adding a third I/O Drawer will require two RIO cables since the third drawer must be on a new loop. If two RIO loops have been configured, adding a third drawer will require one RIO cable. The length of cable selected will depend on whether the third drawer is being installed in the same I/O Rack as an existing drawer.

A fully-configured four drawer system must have two RIO loops. An example of this is shown in Figure 47 on page 188.




Figure 47. Two RIO Loops Attached to Four I/O Drawers

RIO cables are available in three different lengths. The 3-meter cables can only be used to interconnect two I/O Drawers in the same rack. Manufacturing will determine the placement and cabling of I/O Drawers based on the quantity of I/O Racks and RIO cables ordered.

The following remote I/O cables are available on Model S80 systems:

- (# 3142) 3-meter drawer-to-drawer remote I/O cable
- (# 3143) 6-meter rack-to-rack remote I/O cable
- (# 3144) 15-meter rack-to-rack remote I/O cable

5.3.5 Model S80 I/O Bridge Bus

The local I/O bridge bus is a reduced signal version of the system bus that has been optimized for I/O. The I/O bridge bus uses a multiplexed 64-bit address and data path. The I/O bridge bus is parity checked for address, data and control errors. Each bus request is range checked and positively acknowledged for improved error detection. The I/O bridge bus operates in pipeline mode. New requests can be issued before previous requests are completed. The bridges and other chips in the I/O path provide significant queuing.

5.3.6 Model S80 PCI Buses

The PCI bridge chips on both systems convert the I/O bridge bus to PCI. There are 14 PCI-compliant slots running at 33 MHz per I/O Drawer. PCI 2.1 cards are supported. Four PCI bridges are used per I/O planar. One of the PCI bridge chips drives two 64-bit PCI slots. The other three PCI bridges each drive one 64-bit PCI slot and three 32-bit PCI slots. This configuration performance balances the load. Five volts and 3.3 volts are available at the slots. Five volt PCI signaling conventions are used.

The 64-bit PCI slots have a maximum throughput of 266 MBps. The 32-bit PCI slots have a maximum throughput of 133 MBps. It is important to note that no PCI-to-PCI bridges are used in this performance-optimized design. PCI-to-PCI Bridges significantly limit the useful bandwidth of the related PCI slots.

5.3.7 Model S80 Firsts

The Model S80 represents significant new packaging, new components, and new levels of function for the RS/6000 family. It is the first RS/6000 to use the significantly-improved copper technology RS64 III processor and its matched memory controller complex. The Model S80 is the first RS/6000 to support 24-way SMP computing. It is also the first to support 64 GB of memory. The S-Series models are the first RS/6000s to support an expandable I/O subsystem that is packaged separately from the rest of the system.

The S-Series is the first RS/6000 with hardware-based memory scrubbing, which reduces soft error exposures to very low levels. The S-Series is also the first in the RS/6000 family to offer mainframe-class, burnt-in components, which significantly improve reliability over typical industry devices. The S-Series also features new levels of error recovery, detection, and reporting.

The S-Series is the first RS/6000 to use Predictive Failure Analysis to monitor the error rates on DASD, memory, processors, L1 and L2 caches, and remote

I/O. The Model S80 is the first RS/6000 system to allocate one out of every 14 main memory chips for use as a replacement in the event of single-bit errors. Redundancy is also provided for data traffic between the CEC and the I/O Drawer via a looped cable. If the cabling to the I/O Drawer is experiencing errors, the hardware will redirect the I/O traffic through an alternate path on the loop.

5.3.8 Model S80 Configurations

Table 61 lists the S80 minimum configuration.

Table 61. Model S80 Minimum Configuration

Model S80 Minimum Configuration	
Microprocessor	One 6-way 450 MHz RS64 III CPU card
Level 1 (L1) cache	128 KB data/128 KB instruction
Level 2 (L2) cache	8 MB per processor
RAM (minimum)	2 GB
Memory bus width	Quad 512-bit
Ports	One parallel, two serial, one keyboard, and one mouse
Internal disk drive	One 9.1 GB Ultra SCSI (hot-swappable)
Media bays	Two (one available)
Expansion slots	Fourteen PCI (eleven available)
PCI bus width	32- and 64-bit
Memory slots	16
CD-ROM drive	32X (max)
Service processor	Yes
Diskette drive	1.44 MB 3.5-inch diskette drive
SCSI adapters	Two Ultra SCSI PCI adapters
AIX operating system version	Version 4.3.3 unlimited user license

Table 62 lists the Model S80 system expansion capabilities.

Table 62. Model S80 System Expansion

Model S80 Maximum Configuration	
SMP configurations	Up to three additional 6-way processor books
RAM	Up to 64 GB
Internal PCI slots	Up to 56 per system (53 available)
Internal media bays	Up to eight per system (seven available)
Internal disk bays	Up to 48 (hot-swappable)
Internal disk storage	Up to 873.6 GB

5.3.9 Model S80 Additional Features

This section describes the internal features that can be added to a configuration at some additional cost.

The status of a feature is indicative of the following qualifications:

- A** Indicates features that are available and orderable on the specified models.
- S** Indicates a feature that is supported on the new model during a model conversion. These features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed.
- W** Indicates features that are supported but were withdrawn December 13, 1999 or earlier.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, cables, and monitors are not included.

Table 63 lists Model S80 optional features and their status.

Table 63. Model S80 Optional Features

Feature Code	Description	Status
Processors		
5318	RS64 III 6-way 450 MHz first	A
5319	RS64 III 6-way 450 MHz additional	A
Memory		

Feature Code	Description	Status
4190	1024 MB (4 x 256 MB)	A
4191	2048 MB (4 x 512 MB)	A
4192	4096 MB (4 x 1024 MB)	A
4193	8192 MB (4 x 2048 MB)	A
4194	16384 MB (4 x 4096 MB)	A
Host Attachment		
2751	ESCON Control Unit	A
8396	SP Attach Adapter	A
Internal Disk Drives		
2919	9.1GB Ultra-SCSI 16-bit Hot Swap Disk Drive	S
3019	9.1GB F/W Ultra SCSI Disk Drive	S
2901	4.5 GB Ultra SCSI Hot-Swap	S
2911	9.1 GB Ultra SCSI Hot-Swap	S
3002	9.1 GB 10K RPM 1" Ultra SCSI Hot-Swap	A
2913	9.1 GB 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
3117	18.2 GB Ultra 10 K RPM 1" Hot-Swap	A
Internal Tape Drives		
6147	5/10 GB 8 mm	S
6154	20/40GB 8 mm (White)	S
6156	20/40 GB 8 mm (Black)	A
6159	12/24 GB 4 mm	A
Internal CD-ROMs		
2619	20X Speed CD-ROM	S
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2838	GXT120P	S

Feature Code	Description	Status
2830	GXT130P	A
SCSI Adapters		
6204	PCI Universal Differential Ultra SCSI Adapter	A
6206	Ultra SCSI SE	A
6208	SCSI 2 Fast / Wide	S
6209	SCSI2 Fast/Wide Differential	S
6205	Dual Channel Ultra2 SCSI	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	S
6222	SSA Fast-Write Cache Option	A
6235	IBM Advanced Serial RAID Cache Option	A
6230	Advanced Serial RAID Plus Adapter	A
6231	128 MByte DRAM Option Card	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
2948	ARTIC960Hx 4-Port T1/E1 PCI	R
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-ring adapters		
2920	Token-ring adapter	S
2979	Auto LANStreamer token ring	S
4959	Token-ring PCI adapter	A
Ethernet Adapters		

Feature Code	Description	Status
4951	4-Port Ethernet	A
2968	IBM 10/100 Mbps Ethernet	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2986	Fast Etherlink XL 3Com	S
2987	Ethernet AUI / RJ-45	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
Fiber Channel Adapters		
6227	Gigabit Fibre Channel Adapter	A
FDDI Adapters		
2741	SysKonnnect SK-NET FDDI-LP SAS	A
2742	SysKonnnect SK-NET FDDI-LP DAS	A
2743	SysKonnnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A

5.3.10 Model S80 Remote I/O Subsystem Description

The I/O subsystem on both the Model S80 and Model S70 Advanced consists of up to four I/O Drawers housed in I/O Racks and connected to the system CEC. Each I/O Rack can be configured with up to two I/O Drawers, the Primary and Secondary.

5.3.10.1 Model S80 I/O Rack Description

The I/O Rack supplied as a feature of an S-Series server is the 7014-S00 rack. It has 32 EIA units of space and is available with three Power Distribution Unit (PDU) options. They are selected with the following base feature codes:

- (# 9171) Side-Mounted, 1-Phase
- (# 9173) Side-Mounted, 3-Phase
- (# 9174) Side-Mounted, 3-Phase, Swiss

Each PDU provides six power outlets. An optional additional PDU can be configured in a rack providing the option to power additional equipment or providing a degree of high-availability to devices that have multiple redundant power supplies, such as the 7133-D40 SSA enclosure. The following are the feature codes of available additional Power Distribution Units:

- (# 6171) Side-Mounted, 1-Phase
- (# 6173) Side-Mounted, 3-Phase
- (# 6174) Side-Mounted, 3-Phase, Swiss

The following rules have to be followed for a valid configuration:

- A maximum of four I/O Racks can be ordered as features (# 7000) per system order. I/O Drawers should be spread between I/O Racks but not exceed two per rack.
- I/O Racks ordered as feature numbers of the Model S70 Advanced or Model S80 (# 7000) must contain an I/O Drawer. If additional external communication and storage devices, such as 7133 and 7027, do not fit in the space remaining in the I/O Racks, additional empty I/O Racks should be ordered. The additional I/O Racks should be ordered as products (7014-S00 or 7015-R00) rather than features of the S-Series server. There is no limit to the quantity of 7014-S00 and 7015-R00.
- If the quantity of power distribution units (# 6171 or # 6173) is greater than one, the space available in the rack is reduced by one EIA.
- The 3570-B11 and 3570-B12 can only be installed in a specific location in the I/O Rack due to drawer interference. Ordering a machine with a rack contents specify feature code (# 0166 or # 0167) will have the base SCSI I/O Drawer installed in a lower location in the rack, thus leaving a nine EIA space empty at the top of the rack. The 3570 is field installed.
- The 3570-C11 and C12 can be installed in any location within the I/O Rack.
- Many 3490 and 3590 tape libraries cannot be installed in S-Series servers' I/O Racks due to interference with the rack door. These machine types are supported for attachment to the S-Series servers, but they must be in a separate rack (7015-R00) unless the I/O Rack door is removed. Instructions for proper removal of the rack door are included in the tape library installation instructions.
- Two IBM 3590-B11 or 3590-E11 tape drives may share a twelve EIA space in the rack. (Order one placement for each pair of units.)
- The 3490 Model F11 and 3590 Model B1A can be installed in the I/O Rack without interference with the rack door installed.

5.3.10.2 Model S80 I/O Drawer description

The I/O Drawer offers the advantage of fully-redundant power and fans that can be serviced without taking the system down. The drawer also improves the overall cooling with the use of more robust fans, which are especially useful if one of them has failed. In addition to the hot-swappable fans and power supplies, the drawer supports Ultra-SCSI adapters, which are separately cabled to the two disk six packs. The drawer has a local display panel and reports more information for status monitoring. When attached to the Model S80, several PCI adapters are supported in higher capacity configurations than on the Model S70 Advanced.

The primary I/O Drawer contains the I/O planar and the service processor card. In addition, the primary I/O Drawer contains six hot-plug disk bays, one available media bay, one floppy disk drive, one CD-ROM, 14 PCI slots, one keyboard port, one mouse port, two serial ports, and one parallel port. The I/O subsystem is expandable by attaching up to three additional I/O Drawers to a single CEC.

The 14 PCI I/O slots consist of five 64-bit and nine 32-bit PCI slots. Depending on the media and disk configuration chosen, between two and four of the fourteen slots in the first I/O Drawer are used for the service processor, storage, and media support. The remaining slots are available to support graphics, communications, and storage in the initial I/O Drawer configuration.

There is a maximum of two I/O Drawers (# 6320) per I/O Rack. RIO cables and SPCN cables must be ordered for each additional drawer. The manufactured configuration of I/O Drawers in I/O Racks is based on cable lengths ordered.

Figure 48 on page 197 shows a 10 EIA I/O Drawer when viewed from the front. Figure 49 on page 198 shows a 10 EIA I/O Drawer from the rear.




Figure 48. Front View of 10 EIA Drawer

Each drawer has four PCI buses per drawer: Slots 1-4 (PCI bus 1), 5-8 (PCI bus 0), 9-10 (PCI bus 3), and 11-14 (PCI bus 2). Slots 1, 5, 9, 10, and 14 are 64-bit slots. The remaining slots are 32-bit. All supported 32-bit adapters also function in the 64-bit slots. All slots are 33 MHz.

The I/O Drawer has redundant power and cooling. If the I/O Racks are ordered with an additional PDU, each drawer can be connected to two PDUs, thus making the I/O system highly available from a power supply point of view.




Figure 49. Rear View of 10 EIA I/O Drawer

5.3.11 Model S80 Service Tools

The following sections discuss the available tools used for service.

5.3.11.1 Model S80 Service Processor

The service processor provides for excellent RAS service features such as the call home on a hard failure causing a *system down* situation. In addition, the service processor provides functions such as power-on/off of the system, reading the service processor and POST error logs, reading vital product data (VPD), changing the bootlist, viewing boot sequence history, and changing service processor configuration parameters. All of these can be performed remotely. Customers also have the capability to enable console mirroring on the system console so that they can monitor all remote console activity.

5.3.11.2 The Service Director

Service Director generally automatically report problems based on default settings. The customer may modify the default values to prevent the Service Director from placing a service call during hardware upgrades, testing, or in the event that the failed component is not covered by an IBM Service

Agreement (for example, a third-party disk subsystem). The customer can also configure the product to only alert personnel within the customer's IT department (via e-mail). This function can be configured instead of, or in combination with, the ability to automatically place a service call to the IBM Service Center.

5.3.11.3 Dynamic Processor Deallocation

Another high-availability feature is processor de-allocation. This feature allows the system to completely isolate a processor that has been determined to be operating below an acceptable level of reliability. To accomplish this, the system monitors processor status in real-time to determine when the number of recoverable errors has exceeded a specified threshold. When the threshold is exceeded, the system sends a report to the operating environment that a processor is unreliable. The AIX operating environment then re-routes all interrupts and processes to other processors and disables the unreliable processor complex without interrupting end user services. Firmware level 991216 must be installed on the Model S80. Since it is a feature that was introduced to AIX in February 7, 2000, the AIX Section of this book will discuss this function further.

The location of the URL that microcode or firmware can be downloaded from is:

<http://www.austin.ibm.com/support/micro/download.html>

5.3.12 Upgrade Paths to Model S80

Installed Model S70 Advanced systems can be converted to Model S80 systems. This conversion requires replacement of the entire Central Electronic Complex (CEC), including processors and system memory. The support processor group, which includes the service processor card, must also be exchanged to provide the proper firmware for interfacing to the Model S80 CEC. The Model S70 Advanced I/O Rack and drawers, PCI adapter cards, hard disks, and media devices carry forward to the upgraded Model S80 system. This conversion preserves the existing system serial number.

The model upgrade consists of the following items:

- Central Electronics Complex (CEC) - The Model S70 Advanced CEC being replaced is returned to IBM.
- Model S80 labels - Preserving the customer's serial number.
- Model S80 publications.

Processor cards and I/O Drawer attachment hardware must be ordered separately. The following items are *required* and must be ordered separately:

- One - First RS64 III Processor, 6-Way SMP 450 MHz, 8 MB L2 Cache (# 5318)
- 2 GB Minimum System Memory
- Cabling for connection to the primary I/O Drawer includes:
 - One - System Control And Initialization Cable - Upgrade Indicator (# 8006)
 - Two - Power Control Cable - Upgrade Indicator (# 8008)
 - Two - Remote I/O Cables (# 3143 or # 3144)
- One - Remote I/O Hub - Dual Loop - (# 6503)
- One - Support Processor Group - Model Upgrade Only (# 8326)

Note

The S70 Advanced Support Processor being replaced is returned to IBM.

The following items must be carried forward from the system being upgraded:

- I/O Drawer attachment cables (# 8006 and # 8008 above)
- I/O Rack
- Primary and secondary I/O Drawers (including backplanes, SCSI adapters, and cables)
- PCI adapters
- Hard disk drives
- Media devices

5.3.12.1 S70 Advanced to S80 Processor Conversion

Model S70 Advanced systems that are converted to Model S80 systems will require replacement of all system processors. The first Model S70 Advanced 4-way processor card is replaced with the first Model S80 6-way processor card via feature conversion. This processor conversion is only available at the time of initial model upgrade. A maximum of one S70 Advanced 4-way processor card may be converted to one S80 6-way processor card for each system being converted. The existing Model S70 Advanced 4-way processor card being replaced is returned to IBM.

Table 64 shows the processor feature conversions that are available to Model S70 Advanced customers who convert their systems to Model S80 systems.

Table 64. S70 Advanced to S80 Processor Conversion

From	To
(# 0504) RS64 II, 4-way SMP, 262 MHz	(# 5318) RS64 III, 6-way, SMP, 450 MHz
(# 5312) RS64 II, 4-way SMP, 262 MHz	(# 5318) RS64 III, 6-way, SMP, 450 MHz
(# 5314) RS64 II, 4-way SMP, 262 MHz	(# 5318) RS64 III, 6-way, SMP, 450 MHz
(# 5316) RS64 II, 4-way SMP, 262 MHz	(# 5318) RS64 III, 6-way, SMP, 450 MHz
(# 9404) RS64, 4-way SMP, 125 MHz	(# 5318) RS64 III, 6-way, SMP, 450 MHz

5.3.12.2 S70 Advanced to S80 Memory Conversion

Model S70 Advanced systems that are converted to Model S80 systems will require replacement of all system memory. This memory replacement will be implemented by converting existing Model S70 Advanced memory features to Model S80 memory features. The memory conversions are only available at the time of initial model upgrade and are allowed on a one-for-one feature conversion basis. A maximum of four Model S70 Advanced memory features may be converted to Model S80 memory features for each system being converted from S70 Advanced to S80. The existing Model S70 Advanced memory being replaced is returned to IBM.

The following memory feature conversions, shown in Table 65, are available to Model S70 Advanced customers who convert their systems to Model S80 systems.

Table 65. S70 Advanced to S80 Memory Conversion

From	To
1024 GB	
(# 4173) 1024 MB R1 Memory	(# 4190) 1024 MB Memory
(# 4174) 1024 MB R1 Memory Select	(# 4190) 1024 MB Memory
2048 MB	
(# 4171) 512 MB R1 Memory	(# 4191) 2048 MB Memory
(# 9168) Base 512 MB Memory	(# 4191) 2048 MB Memory
(# 4173) 1024 MB R1 Memory	(# 4191) 2048 MB Memory
(# 4174) 1024 MB R1 Memory Select	(# 4191) 2048 MB Memory

From	To
(# 4175) 2048 MB R1 Memory	(# 4191) 2048 MB Memory
(# 4176) 2048 MB R1 Memory Select	(# 4191) 2048 MB Memory
4096 MB	
(# 4171) 512 MB R1 Memory	(# 4192) 4096 MB Memory
(# 9168) Base 512 MB Memory	(# 4192) 4096 MB Memory
(# 4173) 1024 MB R1 Memory	(# 4192) 4096 MB Memory
(# 4174) 1024 MB R1 Memory Select	(# 4192) 4096 MB Memory
(# 4175) 2048 MB R1 Memory	(# 4192) 4096 MB Memory
(# 4176) 2048 MB R1 Memory Select	(# 4192) 4096 MB Memory
(# 4177) 4096 MB R1 Memory	(# 4192) 4096 MB Memory
(# 4178) 4096 MB R1 Memory Select	(# 4192) 4096 MB Memory
8192 MB	
(# 4171) 512 MB R1 Memory	(# 4193) 8192 MB Memory
(# 9168) Base 512 MB Memory	(# 4193) 8192 MB Memory
(# 4173) 1024 MB R1 Memory	(# 4193) 8192 MB Memory
(# 4174) 1024 MB R1 Memory Select	(# 4193) 8192 MB Memory
(# 4175) 2048 MB R1 Memory	(# 4193) 8192 MB Memory
(# 4176) 2048 MB R1 Memory Select	(# 4193) 8192 MB Memory
(# 4177) 4096 MB R1 Memory	(# 4193) 8192 MB Memory
(# 4178) 4096 MB R1 Memory Select	(# 4193) 8192 MB Memory
(# 4179) 8192 MB R1 Memory	(# 4193) 8192 MB Memory
(# 4180) 8192 MB R1 Memory Select	(# 4193) 8192 MB Memory
16384 MB	
(# 4171) 512 MB R1 Memory	(# 4194) 16384 MB Memory
(# 9168) Base 512 MB Memory	(# 4194) 16384 MB Memory
(# 4173) 1024 MB R1 Memory	(# 4194) 16384 MB Memory
(# 4174) 1024 MB R1 Memory Select	(# 4194) 16384 MB Memory

From	To
(# 4175) 2048 MB R1 Memory	(# 4194) 16384 MB Memory
(# 4176) 2048 MB R1 Memory Select	(# 4194) 16384 MB Memory
(# 4177) 4096 MB R1 Memory	(# 4194) 16384 MB Memory
(# 4178) 4096 MB R1 Memory Select	(# 4194) 16384 MB Memory
(# 4179) 8192 MB R1 Memory	(# 4194) 16384 MB Memory
(# 4180) 8192 MB R1 Memory Select	(# 4194) 16384 MB Memory

5.3.13 Model S80 Configuration Notes

This section contains some advice about configuring S-Series machines, both as stand-alone servers and SP-node attachments.

When configuring S-Series Servers, keep the following points in mind:

- In an HACMP configuration, the standard native serial ports on the S-Series machines are not available for heartbeat cabling; therefore, you must install either the 8-port asynchronous adapter (# 2943) with serial-to-serial cable (# 3125) or the 128-port Asynchronous controller (# 2944) with 128-port Asynchronous cable (# 8136), rack-mountable Remote Asynchronous Node (# 8136), RJ-45 to DB-25 converter cable (# 8133), and serial-to-serial cable (# 3125).
- Either an IBM-supported ASCII terminal with an attachment cable or an IBM-supported graphics display with an attachment cable and graphics accelerator is required for initial setup and must be available locally for service. OEM asynchronous terminals are not recommended since they may not transmit the same character sequence and, thus, will not be recognized by the Service Processor.

5.3.13.1 Two S-Series Servers Sharing an I/O Rack

In special circumstances, the S-Series Servers can be ordered without an I/O Rack. This configuration can be a good way of having multiple S-Series Servers without requiring an I/O Rack for every system. This option saves machine room floor space by reducing the footprint of the systems. It also reduces the overall cost of the system.

This configuration is not difficult to achieve, although some special considerations must be taken when ordering.

When placing an order for multiple systems, the first system should be configured as normal. 10 EIA units of space in the I/O Rack should be

reserved for the I/O Drawer of the second system. This is achieved by selecting feature code # 0177 as a rack contents specify feature.

When configuring the second system, there will be a feature change that indicates a rackless system order. The feature is selected in the rack options configuration screen as shown in Figure 50.

The screenshot shows a 'Rack Options Dialog Box' with several sections:

- Racks:** Three options with dropdown menus:
 - 0 7000 I/O Rack
 - 1 7007 Rackless System Order - Plant Install
 - 0 7009 Rackless System Order - Field Install
- Base Power Distribution Units:** Two radio button options:
 - 9171 Side-Mounted, 1-Phase (selected)
 - 9173 Side-Mounted, 3-Phase
- Additional Power Distribution Units:** Two dropdown options:
 - 0 6171 Side-Mounted, 1-Phase
 - 0 6173 Side-Mounted, 3-Phase
- Options:** Two dropdown options:
 - 0 6529 Central Office Feature
 - 0 6095 Linecord - PDU to External Device
- Drawers:** One dropdown option:
 - 1 6320 SCSI I/O Drawer, 10 EIA
- Drawer Groups:** Two dropdown options:
 - 1 6321 Primary I/O Drawer Group
 - 0 6323 Secondary I/O Drawer Group
- SPCN Cables:** Two dropdown options:
 - 0 6006 Drawer to Drawer Power Control Cable
 - 0 6007 Rack to Rack Power Control Cable
- RIO Cables:** Two dropdown options:
 - 0 3126 Remote I/O Cable - Drawer to Drawer
 - 0 3127 Remote I/O Cable - Rack to Rack

At the bottom, there are 'OK' and 'Cancel' buttons.

Figure 50. Rack Options Dialog Box

There are two ways of configuring a rackless system:

- (# 7007) Rackless System Order - Plant Install - This feature allows a system containing a single (primary) I/O Drawer to be ordered without an I/O Rack. The system primary I/O Drawer will be factory installed in the I/O Rack of a companion system that specifies feature code # 0177 to provide the required 10 EIA units rack space and two power outlets. Both systems must be ordered on the same customer order with the same scheduled factory delivery date.
- (# 7009) Rackless System Order - Field Install - This feature allows a system containing a single (primary) I/O Drawer to be ordered without an I/O Rack. The system primary I/O Drawer will be field installed in the I/O Rack of another system that provides the 10 EIA units rack space and two power outlets required for the I/O Drawer.

Note

7017-S80 I/O Drawers are only supported in 7017-S80 or 7017-S7A I/O Racks. Orders containing feature # 7007 are subject to cancellation if a companion system is not ordered with the same factory delivery schedule and feature # 0177 specified.

5.3.14 Model S80 Publications

Table 66 details the publications for the Model S80.

Table 66. Publications Shipped with the Model S80

Order Number	Title
SA38-0558	<i>RS/6000 Enterprise Server S80 Installation and Service Guide</i>
SA38-0557	<i>RS/6000 Enterprise Server S80 User's Guide</i>
SA38-0538	<i>PCI Adapter Placement Reference</i>
SA38-0509	<i>Diagnostic Information for Multiple Bus Systems</i>
SA38-0516	<i>Adapters, Devices, and Cable information for Multiple Bus Systems</i>

5.4 SP External Node Attach and the S-Series

Model S80 and S70 Advanced servers can function as an attached SMP server within the IBM RS/6000 SP environment operating under the control of the IBM Parallel Systems Support Programs for AIX. Up to 16 such systems may be attached in a single configuration. This interconnection can be accomplished utilizing an Ethernet connection and, optionally, using the IBM RS/6000 SP System Attachment Adapter # 8396.

Some I/O adapters available for the S70 Advanced and S80 systems are not supported in the SP environment and must be removed. Refer to *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280, for a list of currently-supported adapters.

IBM Parallel System Support Programs for AIX Version 3.1 (5765-D51), or later, is required for an S-Series server system to function as an attached SMP server within the IBM RS/6000 SP environment.

Each S-Series server system that is to function as an attached SMP server within the IBM RS/6000 SP environment must have a minimum of one

Ethernet adapter. This adapter must be recognized by the system as *en0* and must reside in slot #5 of the primary I/O Drawer.

The RS/6000 SP-attached server must have the latest system and service processor firmware (microcode) installed.

The Control Workstation (CWS) must have sufficient serial port connections and CPU power to support the RS/6000 SP-attachment. The requirement is two RS-232 attachments for each RS/6000 S-Series server that is to be attached. If the CWS does not have sufficient available RS-232 ports, you may have to add one or more 8-port (# 2943) or 128-port (# 2944) Asynchronous adapters, assuming there are enough free adapter slots. If not, a larger capacity CWS will be required.

The 7043 can only be used on SP systems with up to four frames. This limitation applies to the number of frames and not the number of nodes. This number includes expansion frames. The 7043 cannot be used for SP systems with SP-attached servers, such as RS/6000 Enterprise Servers Models S80 or S70 Advanced. Each S80 is considered a frame as well as a node. That would mean that two Model S80s and two regular node frames would be the maximum that the 43p could support.

If the Model S80 is a *node* attached to an SP, the Planning guide states that a Model 140 cannot be the CWS and that the Model F50 must be the CWS. This is because the Model S80 uses more CPU resources than other nodes. This is probably why development stated that 7043 not supported. Typically, the Model F50 is configured with 256 MB minimum and four 9.1GB disks as a minimum for two frame systems. It is recommended that an additional 256 MB is added for the next two SP frames and then up to 1GB for larger systems.

The SP-attached server requires attachment to an RS/6000 SP system with at least one frame containing at least one node and (if switch attached) the SPS (16-port) switch. The 49 inch short *LowBoy* frames with any switch and the 79 inch tall frame with the HPS switch are not supported for SP-attachment. Note, the High Performance Switch hardware is not allowed in any configuration, and the tall frame with the 8 port RS/6000 SP system switch is not allowed in any circumstance. A 10 meter switch cable is provided with the SP-attach order through the RS/6000 SP system configurator.

If this is to be a switchless attachment to the RS/6000 S80 and RS/6000 S70 Advanced servers, then the RS/6000 SP system must also be switchless. There must be available (unused) switch node numbers in your SP System Data Repository even though the switch is not being used. That means an

RS/6000 S80 or RS/6000 S70 Advanced server cannot be attached to a single frame system where all 16 nodes are occupied.

Only one RS/6000 SP System Attachment Adapter is permitted in each SP-attached server. This is consistent with the rules for RS/6000 SP system nodes.

A separately chargeable PSSP 3.1 license must be ordered against each SP-attached server serial number.

If you are configuring a system order for an S-Series machine that is to be SP-attached to an existing SP system, you must also perform an MES order against the SP system to configure the S-Series machine as an external node.

The RS/6000 SP system configurator session for all variations of SP-attach will include the following when adding a SP node attachment:

- (# 9122) Node Attachment Feature - An order for this feature will generate an order for the two 15 Meter RS-232 cables for hardware control and S1TERM connectivity between the CWS and the SP-attached server. It also will generate an order for a 10 meter ground cable. It also traps some data so that the RS/6000 SP system configurator session can keep track of how many nodes (real and logical) are in the system.
- (# 9310) Switch Connection Cable - This feature is required only if the SP-attached server is switch attached. It results in the ordering of one 10 meter switch connection cable. The 10 meter cable is the only supported length at this time.
- (# 9123) Frame Attachment Feature - This feature keeps track of how many frames are currently in your RS/6000 SP system. Since the S-Series SP-attached server is both a logical node and a logical frame in the PSSP code logic, it is important to track this information to avoid exceeding allowable RS/6000 SP system limits for the number of frames.
- (# 9222) Node Attachment Ethernet Bnc Boot Feature - This feature will get a BNC cable to allow RS/6000 SP system Ethernet communications and booting with your SP-attached server, whether switch-attached or not.
- (# 9223) Node Attachment Ethernet Twisted Pair Boot Feature - This feature tracks the choice to incorporate the SP-attached servers as part of an RS/6000 SP system Ethernet Twisted Pair network, but it provides no twisted pair cable. As in the past, the customer is responsible for providing their own twisted pair Ethernet cables.

- (# 5765-D51) PSSP 3.1 Software - This is a feature for each SP-attached server license. The features are as follows:
 - (# 5800) 4 mm tape
 - (# 5801) 8 mm tape
 - (# 5803) CD-ROM

The charge features are different as follows:

- (# 4001) PSSP software on an RS/6000 SP system
- (# 4002) PSSP software on an Externally Attached Server

Note

The RS/6000 SP System Attachment Adapter (# 8396) must always be located in slot #10 of the primary I/O Drawer. No adapters may be installed in slots #9 or #11 of the primary drawer when the SP attachment adapter is installed.

For full details on configuring an S-Series server as a node in an SP system, refer to *RS/6000 SP Planning Volume 1, Hardware and Physical Environment, GA22-7280*.

5.5 High-Availability Cluster Server Solution HA-S70 Advanced

The RS/6000 HA-S70 Advanced Cluster Server (HA-S70 Advanced) is a special package-priced, high-availability solution with options for applications, services, and financing. The HA-S70 Advanced Cluster Server is not a new server model. The HA-S70 Advanced Cluster Server consists of currently available IBM products:

- Model S70 Advanced Enterprise Servers
- 7133 Serial Disk Systems
- HACMP/ES high-availability cluster software
- AIX operating system

5.5.1 Solution HA-S70 Advanced Value Proposition

Here are some of the HA-S70's outstanding characteristics:

- Complete high-availability solution for business-critical and e-business applications. System hardware and software with applications, services, and financing options.

- Package-priced, high-availability hardware and software platform. Priced less than the sum of the individual parts.
- Leadership high-availability cluster solution with competitive price and performance versus UNIX alternatives.
- Robust and mature high-availability hardware, cluster software, and operating system.
- New price/performance point for RS/6000 high-end, enterprise-class, high-availability cluster servers and 7133 SSA storage.
- Supported by thousands of AIX applications and middleware. Ideal solution for transaction processing, database, enterprise resource planning (ERP), and e-business network computing applications.
- High-availability implementation guides and scripts to simplify planning and installation of specific applications.

HA-S70 Advanced is a completely integrated package-priced, high-availability solution framework offering excellent flexibility and growth. Customers requiring performance and capacity greater than the base HA-S70 Advanced 2-node server package can easily scale up by adding more 4-way SMP processors, memory, and I/O for each S70 Advanced Enterprise Server and additional SSA disks and 7133 Serial Disk Systems. Up to eight S70 Advanced Servers, or other RS/6000 servers, can be attached to the HA-S70 Advanced Cluster. The HA-S70 Advanced clusters may also be integrated into an RS/6000 SP cluster, with a total capacity of 32 high-availability cluster nodes. An HA-S70 Advanced Cluster provides the capability to detect and recover from hardware or software-related failures. Providing hardware and software failure detection and recovery is unique to the RS/6000 brand.

5.5.2 Solution HA-S70 Advanced Configurations

An HA-S70 Advanced Cluster Server must include the following minimum configuration as provided in Table 67.

Table 67. HA-S7A Minimum Configuration

HA-S7A Minimum Configuration	
RS/6000 Enterprise Server S70 Advanced	Two
HA Cluster Solution Indicator	One per server
RS64 II Processor, 4-way SMP 262 MHz, Right Hand	One per server
Memory	2 GB per server

HA-S7A Minimum Configuration	
I/O Rack	One per server
Primary I/O Drawer	One per server
LAN Adapters	Two per server
PCI SSA Multi-Initiator/RAID EL Adapter	Two per server
7133-D40 Serial Disk Systems	Two
SSA Copper Cables	Six per 7133-D40
AIX Version 4.3.	One per server
High Availability Cluster Multi-Processing for AIX, Version 4.3 licence (5765-D28)	One per server. Choose between HACMP 4.3 Enhanced Scalability features or HACMP 4.3 Enhanced Scalability Concurrent Resource Manager features.

Note

Feature code 8209 can be used to order up to eight additional 9.1 GB SSA Drive Modules. Feature code 8304 is now withdrawn, but it was used to order up to eight 4.5 GB SSA Drive Modules.

Feature code 8218 can be used to order up to eight additional 18.2 GB SSA Drive Modules.

For high availability, only 7133 Serial Disk Systems and SSA disks are supported; SCSI and Ultra SCSI high-availability (HACMP) configurations, are not supported.

Configuration Note

In an HACMP configuration, the standard native serial ports on the Model S7A or S80 are not supported for heartbeat cabling. Therefore, you must install either the 8-port Asynchronous adapter (# 2943) with a serial-to-serial cable (# 3125) or the 128-port Asynchronous controller (# 2944) with a 128-port Asynchronous cable (# 8136), rack-mountable Remote Asynchronous Node (# 8136), the RJ-45 to DB-25 converter cable (# 8133), and the serial-to-serial Cable (# 3125).

HA-S70 Advanced configuration flexibility allows each S70 Advanced Server to be vertically scaled to a maximum of:

- 12 262 MHz processors
- 32 GB of memory
- 48 9.1 GB Ultra SCSI disks

The 7133 disk system includes eight 9.1 or 18.2 GB disks. Additional 9.1, or 18.2 GB drives can be populated to increase a 7133 drawer to a total capacity of 291.2 GB of SSA storage. For added scalability, up to eight S70 Advanced Servers, up to 32 I/O drawers, and up to 54 TB of 7133 storage may be attached to an HA-S70 Advanced Cluster.

5.5.2.1 HA-S70 Advanced Solution Enhancements

The HA-S70 Advanced Solution Enhancements include the following:

- Two (required) 7133-D40 Serial Disk Systems with eight 9.1 GB or extra-price 18.2 GB SSA disk drives replace the two 7133-020 Serial Disk Systems with eight 9.1 GB SSA disk drives as part of the HA-S70 Advanced packaged-priced solution.
- HACMP 4.3
- A single I/O rack option is available to mount the two required S70 Advanced I/O drawers and the two required 7133-D40 Serial Disk System drawers. Additional I/O racks are also available.
- Optional, additional pretested middleware and applications with HACMP scripts: BEA Systems TUXEDO, Candle, Oracle Flow Manufacturing, PeopleSoft, and TXSeries.
- Extension of the optional special preferred rate financing to qualified customers through December 31, 1999.

5.6 High-Availability Cluster Server Solution HA-S80

The RS/6000 HA-S80 Cluster Server is a high-availability solution with options for applications and services.

5.6.1 Solution HA-S80 Value Proposition

This is a complete high-availability solution for business-critical and e-business applications comprised of system hardware and software with applications and services options.

The HA-S80 Cluster Server is a two-node, high-availability cluster consisting of:

- Two Model S80 Servers
- Two 7133-D40 Serial Disk Systems
- Unlimited AIX Version 4.3.3
- HACMP/ES 4.3. high-availability cluster software
- One or two I/O racks

The HA-S80 represents a new level of price and performance for highly-available, high-end, enterprise-class servers. Redundant hardware and software, coupled with system environmental monitoring and alerting functions to give early power-off warnings, plus local and remote diagnostics and system software designed to detect failures and initiate recovery events, are integrated in an attractive package to minimize and eliminate single points-of-failure.

The HA-S80 solution makes available a flexible and scalable (vertical and horizontal) framework to build a highly-available cluster system tailored to individual customer needs.

5.6.2 Solution HA-S80 Minimum Configuration

Each HA-S80 solution order must include a minimum of the following:

Two - 7017-S80 Servers, each incorporating the following:

- One HA Cluster Solution Indicator (# 0500)
- 4 GB memory minimum
- Two S80 I/O Rack PDUs (Power Distribution Units) for redundancy
- Two 9.1 GB or larger SCSI boot disks mounted in separate 6-pack backplanes within the S80's primary I/O drawer
- One ASYNCH adapter for cluster heartbeat and messaging

Note

Only one serial cable (# 3125 or # 8133) required for the two S80 servers.

- Two LAN Adapters for LAN attachment and backup heartbeat and messaging
- Two PCI SSA Adapters

Note

SSA Fast Write Cache (# 6222) is not allowed.

- One HACMP/ES 4.3.1 or later license must be configured

Note

Additional optional 7017-S80 features may be added as desired.
There is a choice between HA-S80 with Two Racks or One Rack initially when beginning the configuration.

Two 7133-D40 Serial Disk Subsystems, each incorporating the following:

- Four Advanced SSA Disk Modules
- One Manufacturing Integration Code (Rochester or Dublin)
- Six Advanced SSA Cables

Note

Additional optional 7133-D40 features such as disks may be added as desired.
Additional 7133s should be added using the PCRS6000's MES/Upgrade option, not as part of the original configuration.

5.6.3 Solution HA-S80 Highlights

The following section is an outline of the hardware that comes with a standard HA-S80 configuration.

5.6.3.1 Solution HA-S80 High-Availability Hardware

Each HA-S80 server node has comprehensive reliability, availability, and serviceability features.

- Error correcting code (ECC) memory and ECC L2 cache to provide single-bit and double-bit error detection, with single-bit error correction, to minimize potential impact on running applications. Online concurrent diagnostic help enable service technicians to correct potential system malfunctions without interrupting operations. The system is designed to restart automatically by rebooting after an unrecoverable software error, hardware failure, or environmentally-induced stoppage (such as a power interruption).

- For extra reliability, each HA-S80 server node includes redundant cooling fans and redundant AC power supplies that can be hot-swapped. Also, a rack-mounted redundant power distribution unit provides dual 220 volt power and an increased level of availability.
- A service processor has integrated system environmental monitoring and alerting functions (such as AC/DC voltage, fan speed, and temperature sensing) to provide early power-off warnings, and it has facilities for error log analysis and alerts. If potential component failures are detected, the service processor can auto dial-out to a service centre without operator intervention to take preventative maintenance measures. Remote maintenance and diagnostic functions, including console mirroring from a remote site, enable dial-in by a service technician, who can then reboot and restore the system.
- The rack-mounted 7133 Serial Disk System drawer packaged with the HA-S80 is designed for high availability. With the 7133's redundant loop design, a single cable failure will not cause loss of access to data. If there is a failure on a loop, the server SSA disk adapter is designed to continue accessing automatically the device in a non-loop configuration. Once the path is restored, the adapter is designed to reconfigure automatically to resume the normal mode of operation. Each SSA disk drive is individually mounted in a docking carrier for easy replacement. If there is a disk failure, the user-replaceable disks that can be hot-swapped can be removed without loss of communication between the adapter and the other disks on the loop. Similar automatic reconfigurations will reset to normal once the disk is repaired or replaced. To further enhance high availability, the 7133 includes redundant power supplies and fans that can be replaced without affecting HA-S80 system operations.

5.6.3.2 Solution HA-S80 High-Availability Software

HACMP is a cluster software program that provides highly-available application services for the HA-S80 and all RS/6000 servers (uniprocessors, SMPs, or nodes of an SP(TM) server). HACMP detects system hardware and application failures. Once a failure is detected, HACMP initiates graceful failover to a backup communications adapter, backup disk adapter, backup power supply, or to a backup recovery server with a minimal loss of time.

One method of delivering high-availability services is the use of redundant disks. HACMP provides a higher level of availability than Redundant Array of Independent Disk (RAID) subsystems or disk mirroring because, in addition to disk failures, it helps protect against network, adapter, system and

The AIX operating system provides significant functions for high availability, such as the following.

- Logical volume (disk) mirroring (LVM), disk controller duplexing.
- Journalled File System (JFS) to maintain file system consistency and prevent data loss.

This clustered approach, together with the capability of applications failover and recovery/restart of the HACMP configured server, provides additional levels of high-availability processing for business-critical applications.

Supported Application:

HA-S80 hardware and software provide a highly-available server framework for a variety of cross-industry and industry-specific application solutions. Industry-leading solution providers, both IBM and non-IBM, provide applications and claim support for RS/6000 servers and RS/6000 high-availability solutions (HA50, HA-S70 Advanced, HA-H70, and HA-S80) that support AIX Version 4.3 or later and HACMP 4.3 or later.

Pretested HACMP scripts have been developed and are available for the following applications and middleware:

- Collaborative
 - Lotus Domino Server
- Database
 - IBM DB2 Universal Database
 - Oracle Database
- Enterprise Resource Planning
 - BAAN applications
 - Oracle Financial applications
 - Oracle Flow Manufacturing
 - SAP R/3 applications
- Enterprise Availability Management
 - Candle Command Centre for High-Availability Systems
 - Tivoli Enterprise
- Transaction Monitor
 - IBM TXSeries in Websphere Application Server Enterprise Edition
 - BEA Tuxedo

The application scripts are executable start/stop failover scripts and are tested for compatibility with the specific application. Implementation guides describe the scripts and provide assistance for planning and installation. The scripts and implementation guides are available *as is* from a central IBM repository and can be downloaded via the Internet by IGS, Business Partner, or customer service personnel from the IBM Web site at:

<http://www.rs6000.ibm.com/HAScripts>

Since this is a dynamic Web site, the current application script releases that are available are specified on the site.

Scalable growth

HACMP software is scalable to address a broad range of business-critical, high-availability application needs. With the HA-S80, up to eight S80 servers can be configured for high availability. HACMP offers built-in growth and helps protect investments in cluster hardware, software, services, and training.

5.6.4 Solution HA-S80 System Expansion

For high availability, only 7133 Serial Disk Systems and SSA disks are supported; SCSI and Ultra SCSI high-availability (HACMP) configurations are not supported.

HA-S80 configuration flexibility allows each S80 Server to be vertically scaled to a maximum of:

- 24 450 MHz processors
- 64 GB of memory
- 48 Ultra SCSI disks

The 7133 disk system must include four disks. Additional drives can be populated to increase a 7133 drawer to a total capacity of 582 GB of SSA storage. For added scalability, up to eight S80 servers, up to 32 I/O drawers, and up to 108 TB of 7133 storage may be attached to an HA-S80 Cluster.

5.7 System Combinations in the S80 I/O Rack

Figure 51 shows that you can have two S80 CECs sharing an I/O Rack.




Figure 51. Two Systems Sharing One I/O Rack

Note

If power is to be removed from one I/O drawer of a system that shares an I/O rack with another system, ensure that the power is removed from the correct drawer. Check the cabling before removing the power to make sure that you do not disrupt the operation of the second system.

The I/O rack can also have other system drawers installed. The following figures show configurations that allow system drawers to be installed in the I/O rack along with primary and secondary I/O drawers.

Figure 52 shows the basic configuration with H50/H70 or H80 sharing space in the S80 I/O Rack.




Figure 52. Basic Configuration with H50, H70, or H80 Sharing the S80 I/O Rack

Note

Ensure that the cables from the S80 or S80 I/O Drawer are tied to the sides of the rack to enable servicing of the system installed under the I/O Drawers.

Figure 53 shows a system with one I/O Drawer (Primary drawer is in top position) and two Model 7026 H-series.




Figure 53. System with One I/O Drawer, Two H50s, H70s, or H80s

Figure 54 shows how to install two I/O drawers and one 7026 H-Series machine in an I/O Rack of the Model S80.




Figure 54. System with Two I/O Drawers with H50, H70, or H80

Figure 55 shows a combination of primary and secondary I/O Drawers in a Model S80 I/O Rack.




Figure 55. Primary and Secondary I/O Drawers in Model S80 I/O Rack

Chapter 6. The RS/6000 Server and Solution Racks

This chapter takes a look at the following models of racks that have housed the RS/6000 and its solutions.

Housing servers, disk storage, or tape storage in one rack makes access and servicing more efficient.

In this chapter, the following models will be outlined:

- Model 7015-R00
- Model 7014-S00
- Model 7014-T42
- Model 7014-T00

6.1 Model 7015-R00 Overview

Since May 24, 1994, the Model 7015-R00 System Rack provides mounting space and support for the R-Series uniprocessor and symmetric multiprocessor (SMP) drawers as well as external supported devices. The 1.6 meter high rack meets the requirements of the Electronics Industries Standards (EIA) IA-310C standard. An EIA is a unit measure of vertical mounting space that equals 44.45 mm (1.75 inches). The R00 System Rack has 32 EIA units of vertical mounting space.

The enclosure accepts 19-inch panels and mounting rails with a universal spacing EIA hole pattern. The Model R00 enclosure is designed to occupy minimal floor space. Dimensions are 650 mm wide (25.6 inches) by 940 mm deep (37.0 inches), for a total floor space requirement of 0.6 square meters (6.4 square feet).

The RS/6000 Model R00 System Rack enclosure has solid side and bottom panels. The top is solid except for three cable access holes. The front is open for installation of electronic equipment, and the rear is covered by a door that may be removed if installation access requires it. Cables may enter and exit the Model R00 System Rack either at the top or in the rear at floor level. Several options are available to provide either AC or DC power distribution for the drawers/enclosures mounted in the rack.

The Model R00 System Rack is mounted on casters to make it easy to move. Once in place, the casters can be secured for increased stability. Optional

features are available to provide additional capability to bolt the rack to either raised or concrete floors.

On all initial orders, the factory installs blank filler panels in unused rack space and bypass plugs in unused outlets. The customer does not need to order these features.

6.1.1 Model 7015-R00 Physical Specifications

Table 68 outlines the physical size and weight specifications of the model 7015-R00.

Table 68. Model 7015-R00 Net Weight and Physical Size

Dimension	Size
Width	650 mm (25.5 inches)
Depth	940 mm (37.0 inches)
Depth with SMP ¹	1057 mm (41.6 inches)
Height	1.6 meters or 1578 mm (62 inches)
Empty Weight	130 kg (286 pounds)
Maximum Weight	435 kg (960 pounds)
¹ SMP depth includes deep door (# 6096)	

6.1.1.1 Model 7015-R00 Power System

The R00 System Rack power distribution system provides 200 to 240 Volts AC (50 Hz/60 Hz) to the system components (drawers). The distribution system is capable of providing up to 4.8 KVA per line cord. A maximum of two power distribution units (line cords) may be installed. Feature codes # 6110 and # 6115, which were withdrawn May 8, 2000, replace the AC input Power Distribution Unit with a DC input Power Distribution Panel for operation with -48V power.

The R00 System Rack may be used as the base rack to support one or more CPUs with available external supported devices as well as an expansion rack to support additional external supported devices. AC and DC power distribution units, with or without battery backup for AC, are available. Features are available to provide bolt-down to either concrete or raised floor.

The Model 7015-R00 is ivory in color.

6.1.2 Model 7015-R00 Configuration Suggestions

Consider the following suggestions when configuring an R00 rack:

- A maximum of 32 EIA units is available for mounting drawers.
- A maximum of six drawers may be installed with a single Power Distribution Unit (PDU).
- More than six drawers may be installed as space is available if the optional additional PDU (# 6171 or # 6173) is installed.
- A maximum of four processor drawers may be ordered in one, multisystem order.
- If one or more Model R00 racks is ordered, then all drawers in that multisystem order must be mounted in those racks. (Unmounted drawers must have separate multisystem orders.)
- Two IBM 3490 Model E11s may share an 8 EIA space in the rack. (Order one placement code for each pair of units.)
- Two IBM 3590 Model B11s may share a 12 EIA space in the rack. (Order one placement code for each pair of units.)
- Two IBM 3590 Model B11s or E11s may share a 12 EIA space in the rack (Order one placement code for each pair of units).
- If any one of the *bolt-down* features (# 6090, # 6091, # 6092, # 6093, # 6190, # 6192) is ordered, then only two SMP processor drawers may be installed in that rack.
- A total of three side-mount power-distribution units (# 6171, # 6173, # 6174, # 9171, # 9173, # 9174) may be installed in one rack. The third one will require one EIA of space at the bottom of the rack.
- Each R00 rack with AC power must have one (only) of the following features: # 6091, # 6190, or # 9094. # 9094 should be ordered unless a special Telco installation is required.
- Each R00 rack with DC power must have one (only) of the following features: # 6092, # 6093, or # 6192, plus either # 6110 or # 6115.
- AC and DC power cannot be mixed within a multisystem rack order.
- One of the following features must be ordered for each # 9910 UPS: # 9171, # 9173, # 6171, or # 6173.
- Feature # 6096, deep door for rack, must be ordered if an SMP drawer is installed.
- Refer to the Power Cord section to determine the correct type of Power Distribution Unit for the country where the system will be installed.

6.1.3 Model 7015-R00 Publications

The following publications, shown in Table 69, are shipped with the Model 7015-R00

Table 69. Publications Shipped with the Model R00

Order Number	Title
SA23-2628	<i>7015 Installation and Service Guide</i>
SA23-2652	<i>System Unit Safety Information</i>

6.2 Model 7014-S00 Overview

The Model 7014-S00 was announced on October 6, 1997. The IBM 7014 Model S00 is a 1.6-meter-high rack for mounting RS/6000 Model S80 I/O drawers or other rack-mountable external devices. The rack meets the Electronics Industries Standards (EIA) 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 rear-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 stabilizing the rack for central office environments. Figure 56 on page 227 shows the Model 7014-S00 Rack front and rear views.




Figure 56. Model 7014 S00 Rack-Front (Door Removed) and Rear

6.2.1 Model 7014-S00 Physical Specifications

Table 70 outlines the physical size and weight specifications of the Model 7014-S00.

Table 70. Model 7014-S00 Net Weight and Physical Size

Dimension	Size
Width	650 mm (25.5 inches)
Depth	1019 mm (40.1 inches)
Height	1.6 meters or 1578 mm (62 inches)
Empty Weight	130 kg (286 pounds)
Maximum Weight	435 kg (960 pounds)

6.2.1.1 Model 7014-S00 Power System

The S00 System Rack power distribution system provides 200 to 240 VAC (50Hz/60Hz) to the system components (drawers). The distribution system is capable of providing up to 4.8 KVA. The optional DC Power Distribution Panels (# 6110 and # 6115) replace the AC input Power Distribution Unit with a DC input Power Distribution Panel.

With AC Power Distribution Unit:

Specify one of these features numbers in the S00 System Rack order if you plan to use it with the standard AC power distribution unit:

- # 9171 Power Distribution Unit - Base, Single Phase
- # 9173 Power Distribution Unit - Base, Two of Three Phase
- # 9174 Power Distribution Unit - Base, Three Phase - Swiss

With DC Power Distribution Unit:

Installations that require minus 48 Volts DC input power may select one of the following power distribution options. Individual drawers must also be ordered with the minus 48 VDC option:

- # 6110 DC Power Distribution Select for CPU Rack
- # 6112 Dual DC Power Distribution Select for I/O drawer
- # 6115 DC Power Distribution Select for non-CPU Rack

The Model 7014-S00 is black in color.

6.2.2 Model 7014-S00 Configuration Suggestions

The following are suggestions for when you configure an S00 rack:

- A maximum of 32 EIA units is available for mounting drawers.
- A 7015-R00 rack may also be interchangeably used. However, it must be noted that the S00 is black, and the R00 rack is painted white and has no door.
- A maximum of six drawers may be installed with a single Power Distribution Unit (PDU).
- Up to four Model H50/H70 drawers, depending on their configurations, may be installed in the IBM 7014 RS/6000 Rack Model S00.
- More than six drawers may be installed as space is available if the optional additional PDU (# 6171), (# 6173), or (# 6174) is installed.
- A maximum of four processor drawers may be ordered in one, multisystem order.
- If one or more Model S00 racks is ordered, then all drawers in that multisystem order must be mounted in those racks. (Unmounted drawers must have separate multisystem orders.)
- Two IBM 3490 Model E11s may share an 8 EIA space in the rack. (Order one placement code for each pair of units.)

Note

Many 3490 and 3590 tape libraries cannot be installed in the 7014-S00 rack due to interference with the rack door. Instructions for proper removal of the rack door are included in the tape library installation instructions.

- Two IBM 3590 Model B11s may share a 12 EIA space in the rack. (Order one placement code for each pair of units.)
- AC and DC power cannot be mixed within a multisystem rack order.
- Central Office Feature (# 6529) is a prerequisite for DC Power Distribution Panel (# 6110, # 6112, or # 6115).
- A total of three side-mount Power Distribution Units (# 6171, # 6174, # 9171, # 9173, or # 9174) may be installed in one rack. The third PDU will require the addition of 1 EIA space at the bottom of the rack.
- Uninterruptible Power Supply 9910-U33) requires one of the following features (# 6171, # 6173, # 9171, or # 9173) for each # 9910 installed.
- Refer to the Power Cord section to determine the correct type of Power Distribution Unit for the country where the system will be installed.
- (# 0500) - HA Solution Indicator with Redundant Power: This feature may be required when ordering certain High Availability Solution offerings. It identifies the system as part of an HA Solution order and provides the appropriate redundant power features. For the 7014-S00 rack, # 0500 provides second Power Distribution unit equivalent to (# 6171), (# 6173), or (# 6174). For the 7026-H50 server, # 0500 provides a Redundant AC Power Supply equivalent to # 6296.

Drawer limits and possible drawer combinations are as follows:

- Two IBM Model S7A or S80 I/O drawers.
- Four H50s or four H70s or a combination of four is allowed in the S00.
- Two IBM 3490-C11s or C22s: No other drawers allowed.
- Two IBM 3490-C11s or C22s: No other drawers allowed.

6.2.3 Model 7014-S00 Publications

Table 71 lists the publications that are available for the Model 7014-S00

Table 71. Publications Available for the Model 7014-S00

Order Number	Title
SA38-0550	7014 Model S00 Rack Install and Service Guide

Order Number	Title
SA23-2652	<i>System Unit Safety Information</i>

6.3 Model 7014-T00 Overview

The Model 7014-T00 was announced on May 9, 2000.

The 1.8 meter 7014-T00 rack has 36 EIA units (36U) of vertical mounting space for 19-inch mounting-width drawers. An EIA unit is an industry standard for indicating vertical mounting space in a rack, where 1U equals 44.45 mm (1.75 inches). This rack meets the Electronics Industries Association EIA-310-D standard. The standard color is black (Figure 57), which is consistent with the Model S80 Enterprise Server and other recent RS/6000 models. The racks are also orderable in a white color that matches the Model R00 rack and other older RS/6000 models.



Figure 57. Model 7014-T00 Front View

6.3.1 Model 7014-T00 Highlights

The 1.8 meter Model T00 is the rack that will best satisfy most general purpose requirements. It is compatible with past and present RS/6000 racks and is designed for use in all situations that have previously used the older rack Models R00 and S00. New and improved features include:

- 36 EIA units (36U) of usable space
- Optional removable side panels
- Optional front door
- Optional side-to-side mounting hardware for joining multiple racks
- Increased power distribution and weight capacity
- Improved cabling with centered drawers and integrated cable strain relief
- Standard black or optional white color
- Taller, stronger 19-inch wide rack provides increased height, weight, and power distribution capacity for RS/6000 systems
- Optional, ruggedized rack feature provides added earthquake protection with modular rear brace, concrete floor bolt-down hardware, and bolt-in steel front filler panels
- Model T00 supports both AC and DC configurations

6.3.1.1 Side Panels and Side-to-Side Rack Suite Attachment

Unlike previous RS/6000 racks, the rack side panels:

- Are optional
- Can be quickly installed or removed from the outside the rack
- Require no tools

This design allows a row of racks without side panels to be bolted together in a continuous suite through use of the optional side-to-side rack connecting hardware. When multiple racks are joined in this way, cables can be run between racks without having to exit the continuous rack enclosure. A small gap is maintained between the two adjacent racks, which is filled by three matching steel trim pieces that:

- Snap into place on the front, top, and rear, between each rack
- Cover the space between each rack

The removable side panels are:

- Needed only for the two end racks of the suite

- Used for the T00 and the T42 racks

On all T42 racks, small side panels are included to cover the portion of the rack that extends above 1.8 meters.

6.3.1.2 Front Door and Front Trim Kit

The front door is optional. The main surface is steel with a perforated hole pattern to provide the required ventilation and sufficient visibility of indicator lights for equipment installed inside the rack.

Note

Each rack must be ordered with either a front door or a front trim kit (but cannot be ordered with both).

The front trim kit consists of steel parts that:

- Snap into the front of the rack at the left, right, and top edges

Because the T42 rack is taller than the T00 rack, the following require different sets of features for the two models:

- Front doors
- Front trim kits
- Side-to-side rack connecting kits

6.3.1.3 Ruggedized Rack Feature

The optional ruggedized rack feature provides additional hardware that:

- Reinforces the rack and anchors it to the floor
- Includes a large steel brace that:
 - Bolts into the rear of the rack
 - Is hinged on the left side so that it can swing out of the way for access to the rack drawers when necessary

The Ruggedized Rack Feature also includes:

- Hardware for bolting the rack to a concrete floor
- Bolt-in steel filler panels for any unoccupied spaces in the rack

6.3.1.4 Other Standard Equipment and Optional Features

The standard equipment shipped with each rack includes:

- A rear door

- Front and rear anti-tilt plates
- Four leveling feet
- Snap-in filler panels for any unoccupied rack space
- An AC power distribution unit (PDU) mounted inside the rack in an area separate from the drawer mounting area
- There is room for up to four PDUs without using any of the rack's 36U or 42U capacity.

The rack drawer mounting area:

- Is centered relative to the two sides of the rack
- Provides room to route cables within the frame on the left and right sides

There are numerous anchor points formed into the frame where cables can be tied down for strain relief. For Telco applications, -48 volt DC power distribution can be ordered instead of the standard AC PDU.

6.3.2 Model 7014-T00 Physical Specifications

The following section outlines the Model 7014-T00 physical specifications.

6.3.2.1 Net Weight and Physical Size

Table 72 outlines the physical size and weight specifications of the model 7015-T00.

Table 72. Model 7014-T00 Net Weight and Physical Size

Dimension	Size
Width with side panels	644 mm (25.4 inches)
Width without side panels	623 mm (24.5 inches)
Depth with front and rear doors	1098 mm (43.3 inches)
Depth with rear door only	1042 mm (41.0 inches)
Height	1804 mm (71 inches) with standard AC power
	1926 mm (75.8 in) with -48 volt DC power
Empty Weight	244 kg (535 lb)
Maximum Weight	816 kg (1795 lb)

6.3.2.2 Model 7014-T00 Power System

The operating voltage may be 200 to 240 V AC 50/60 Hz Power Source Loading: 4.8 KVA per Power Distribution Unit (PDU) (rack may have up to 4 PDUs)

Optional DC system available on Model T00 only: -48 V DC.

The rack power distribution system provides 200 to 240 V AC (50Hz/60Hz) to the system components (drawers). Each PDU is capable of providing 1 to 4.8 KVA. The optional DC Power Distribution Panels (# 6115, # 6116, and # 6117) replace the standard AC input Power Distribution Unit with a DC input Power Distribution Panel.

AC Power Distribution Unit

Specify one of the following features numbers on the T00 or T42 system rack order if you plan to use it with the standard AC power distribution unit:

- # 9171 Power Distribution Unit Base, Single Phase
- # 9173 Power Distribution Unit Base, Two of Three Phase

DC Power Distribution Unit

Installations that require -48 volts DC input power can select one of the following power distribution options. Individual drawers must also be ordered with the -48 volts DC option:

- # 6115 DC Power Distribution Select for Storage Rack
- # 6116 DC Power Distribution Select for Multipurpose Rack
- # 6117 DC Power Distribution Select, Bulkhead Connectors

6.3.3 Model 7014-T00 Configuration Suggestions

The following are suggestions for the configuration of a T00 rack:

- Optional features include a front door and removable side panels
- The Model 7014-T00 is available as a feature (# 7036) on the Model 7017-S80.
- A maximum of 36 EIA units are available for mounting drawers in the rack.
- A DC Power Distribution Panel (PDP), # 6115, # 6116, or # 6117, is available for systems with drawers that have -48 volt DC Input Power features. A rack cannot have both AC and DC power distribution. The PDP is mounted on top of the rack.
- Racks ordered with -48 volts DC power are available only in white.

Suggested rack content can be:

- 7026-H70 - 8 EIA
- 7026-H80 - 5 EIA (maximum of 3)
- 7026-M80- 8 EIA (maximum of 2)
- 3590-B11 - (2 x B11 = 12 EIA)
- UPS - 4 EIA
- 7133-D40 - 4 EIA (maximum of 9)
- 7337-306 - 5 EIA
- 2104-DL1 - 3 EIA
- 3490-F11 - 4 EIA
- 3570-C11 - 6 EIA
- 3590-E11 - (1x OR 2 x= 12 EIA)
- 3570-C12 - 6 EIA

Feature # 6117 DC power distribution select, bulkhead connectors provide a top-mounted, dual DC Power Distribution Panel for a rack containing varying quantities of CPU drawers or storage subsystems. Up to two DC H80 systems or two DC M80 systems are supported in addition to up to four DC storage subsystems. Unlike previous PDPs, such as # 6116, this feature is built without attached power cables. Instead, the # 6117 PDP comes with a series of power connectors built into its rear bulkhead. The appropriate DC power cables are included with the supported drawer systems and plug into the power connectors at the rear of the # 6117 PDP.

6.3.4 Model 7014-T00 Publications

The publication listed in Table 73 are available for the Model 7014-T00.

Table 73. Publications Available for the Model 7014-T00

Order Number	Title
SA38-0577	<i>Model T00 and T42 Rack Installation and Service Guide</i>

6.4 Model 7014-T42 Overview

The Model 7014-T42 was announced on May 9, 2000. The 2.0 meter rack addresses the special requirements of customers who want a very tall enclosure to house the maximum equipment in the smallest possible floor space. It may require special shipping and handling procedures because of

its height. It provides 42 EIA units (42U) of usable space in addition to all the features available in the T00. An EIA unit is an industry standard for indicating vertical mounting space in a rack, where 1U equals 44.45 mm (1.75 in). These racks meet the Electronics Industries Association EIA-310-D standard. The standard color is black, which is consistent with the Model S80 Enterprise Server and other recent RS/6000 models. The racks can also be ordered in a white color that matches the Model R00 rack and other, older RS/6000 models.

6.4.1 Model 7014-T42 Highlights

The 2.0 meter Model T42 rack highlights are:

- It has a capacity of 42U.
- Addresses the special requirements of customers who want a very tall enclosure to house the maximum equipment in the smallest possible floor space.
- May require special shipping and handling procedures because of its height.
- Improved cabling with centred 19-inch wide drawers and integrated cable strain relief.
- The Model 7014-T42 supports AC only.
- Small side panels on the T42 are always included to cover the portion of the rack that extends above 1.8 meters.
- Optional features include a front door and removable side panels.
- Each rack must be ordered with either a front door or a trim kit. A rack cannot have both a front door and a trim kit.
- A rack must not have a side panel installed where it is joined to an adjacent rack with the rack suite attachment kit. A suite of racks with a quantity of N racks joined together would typically require a quantity of N-1 rack suite attachment kits and two side panels to cover the two ends of the suite of racks.
- Rack suite attachment kits are intended for joining two racks of equal height. Joining a T00 rack to a T42 rack is not recommended. If this is done, the bolts and spacers will function correctly, but the trim pieces are not designed for use with uneven rack heights.
- Front doors are not available on white racks.
- There are six AC power outlets on each Power Distribution Unit (PDU). A maximum of six drawers can be installed in a rack that has a single PDU.

- More than six drawers can be installed when space is available in the rack if additional PDUs (# 6171, # 6173, or # 6174) are installed.
- More than six drawers can be installed when space is available in the rack if additional PDUs (# 6171 or # 6173) are installed.
- If one or more racks are ordered, all drawers in that multisystem order must be mounted in those racks. (Unmounted drawers must have separate multisystem orders.)
- Two IBM 3490 Model F11s can share an 8-EIA space in the rack (order one placement code for each pair of units).
- Two IBM 3590 Model B11s or E11s may share a 12-EIA space in the rack (order one placement code for each pair of units).
- Many 3490 and 3590 tape libraries may interfere with the front door of a T00 or T42 rack. When planning for the installation of this type of equipment, racks should be ordered without a front door.
- The available Uninterruptible Power Supply (9910-U33) requires one feature (# 6171, # 6173, # 9171, or # 9173) for each 9910-U33 installed.
- The factory will install drawers only in the lower 32-EIA units of the Model T42 rack. Installation of drawers in the upper 10-EIA units must be done at the customer's location.
- The Model 7014-T42 is available as a feature (# 7037) on the Model 7017-S80.

Not only is the Model 7014-T42 taller, but it is stronger with 19-inch wide racks. Optional features include the front door and the side panels. This allows for side-to-side joining of multiple racks into suites. The optional ruggedized rack feature provides added earthquake protection with:

- Modular rear brace
- Concrete floor bolt-down hardware
- Bolt-in steel front filler panels

The racks are available in standard black or optional white color.

The Model 7014-T42 supports AC only.

6.4.1.1 Side Panels and Side-to-Side Rack Suite Attachment

Unlike previous RS/6000 racks, the rack side panels:

- Are optional
- Can be quickly installed or removed from the outside the rack

- Require no tools

This design allows a row of racks without side panels to be bolted together in a continuous suite through use of the optional side-to-side rack connecting hardware. When multiple racks are joined in this way, cables can be run between racks without having to exit the continuous rack enclosure. A small gap is maintained between the two adjacent racks, which is filled by three matching steel trim pieces that:

- Snap into place on the front, top, and rear, between each rack
- Cover the space between each rack

The removable side panels are:

- Needed only for the two end racks of the suite
- Used for the T00 and the T42 racks

On all T42 racks, small side panels are included to cover the portion of the rack that extends above 1.8 meters.

6.4.1.2 Front Door and Front Trim Kit

The front door is optional. The main surface is steel with a perforated hole pattern to provide the required ventilation and sufficient visibility of indicator lights for equipment installed inside the rack.

Note

Each rack must be ordered with either a front door or a front trim kit (but cannot be ordered with both).

The front trim kit consists of steel parts that:

- Snap into the front of the rack at the left, right, and top edges

Because the T42 rack is taller than the T00 rack, the following require different sets of features for the two models:

- Front doors
- Front trim kits
- Side-to-side rack connecting kits

6.4.1.3 Ruggedized Rack Feature

The optional ruggedized rack feature provides additional hardware that:

- Reinforces the rack and anchors it to the floor

- Includes a large steel brace that:
 - Bolts into the rear of the rack
 - Is hinged on the left side so that it can swing out of the way for access to the rack drawers when necessary

The ruggedized rack feature also includes:

- Hardware for bolting the rack to a concrete floor
- Bolt-in steel filler panels for any unoccupied spaces in the rack

6.4.1.4 Other Standard Equipment and Optional Features

The standard equipment shipped with each rack includes:

- A rear door.
- Front and rear anti-tilt plates.
- Four leveling feet.
- Snap-in filler panels for any unoccupied rack space.
- An AC power distribution unit (PDU) mounted inside the rack in an area separate from the drawer mounting area.
- There is room for up to four PDUs without using any of the rack's 36U or 42U capacity.

The rack drawer mounting area:

- Is centered relative to the two sides of the rack
- Provides room to route cables within the frame on the left and right sides

There are numerous anchor points formed into the frame where cables can be tied down for strain relief. For Telco applications, -48 volt DC power distribution can be ordered instead of the standard AC PDU.

6.4.2 Model 7014-T42 Physical Specifications

Table 74 outlines the physical size and weight specifications of the model 7015-T42.

Table 74. Model 7014-T42 Net Weight and Physical Size

Dimension	Size
Width - with side panels	644 mm (25.4 inches)
Width - without side panels	623 mm (24.5 inches)
Depth - with front and rear doors	1098 mm (43.3 inches)

Dimension	Size
Depth - with rear door only	1042 mm (41.0 inches)
Height	2015 mm (79.3 in) with standard AC power
Empty Weight:	261 kg (575 lb)
Maximum Weight	930 kg (2045 lb)

6.4.2.1 Model 7014-T42 Power System

The systems operating voltage is 200 to 240 V AC 50/60 Hz Power Source Loading: 4.8 KVA per Power Distribution Unit (PDU) (rack may have up to four PDUs).

Optional DC system available on Model T00 only: -48 V DC

The rack power distribution system provides 200 to 240 V AC (50Hz/60Hz) to the system components (drawers). Each PDU is capable of providing up to 4.8 KVA. The optional DC Power Distribution Panels, #6115, #6116,1 and #6117, replace the standard AC input Power Distribution Unit with a DC input Power Distribution Panel.

Specify one of these features numbers on the T00 or T42 system rack order if you plan to use it with the standard AC power distribution unit:

- # 9171 Power Distribution Unit - Base, Single Phase
- # 9173 Power Distribution Unit - Base, Two of Three Phase

6.4.3 Model 7014-T42 Configuration Suggestions

A maximum of 42 EIA units are available for mounting drawers in the rack.

The suggested rack content can be:

- 7026-H70 - 8 EIA
- 7026-H80 - 5 EIA (maximum of 3)
- 7026-M80 - 8 EIA (maximum of 2)
- 3590-B11 - (2 x B11 = 12 EIA)
- UPS - 4 EIA
- 7133-D40 - 4 EIA (maximum of 10)
- 7337-306 - 5 EIA
- 2104-DL1 - 3 EIA

- 3490-F11 - 4 EIA
- 3570-C11 - 6 EIA
- 3590-E11 - (1 x or 2 x= 12 EIA)
- 3570-C12 - 6 EIA

6.4.4 Model 7014-T42 Publications

The publication in Table 75 are available for the Model 7014-T42.

Table 75. Publications Available for the Model 7014-T42

Order Number	Title
SA38-0577	<i>Model T00 and T42 Rack Installation and Service Guide</i>

Chapter 7. Large Scale Servers - RS/6000 SP Systems

This chapter provides information about IBM's RS/6000 large scale servers. The servers that fall into this category are called the RS/6000 SP (Scalable POWERparallel) systems. Figure 58 shows an RS/6000 SP Model 550 in a tall frame.



Figure 58. The RS/6000 SP Model 550

7.1 SP Origins

In 1990, IBM's Advanced Workstation Division in Austin, Texas, introduced the RISC System/6000 (RS/6000) family of UNIX-based workstations and servers. These early RS/6000 machines boasted stellar floating point performance for their time, owing to the strength of the Performance Optimization with Enhanced RISC (POWER) CPU architecture.

The High Performance Supercomputer Systems Development Laboratory (HPSSDL) that was formed within IBM's Large Systems Division in Kingston and Poughkeepsie, New York, intended to create a supercomputer based on existing, widespread technology.

The fact that UNIX was entrenched in their target market - large scientific and technical customers - HPSSDL, which was considering mainframe processor technology, experimented with off-the-shelf RS/6000 machines by adding ESCON adapters and interconnecting them with an ESCON Director. The RS/6000 machines were repackaged as nodes and placed in frames. Only five of the large, sheet-metal drawers for the nodes could be placed in one frame. Thus, the predecessor SP was developed.

As HPSSDL was developing smaller drawers, an IBM research group in Yorktown, New York, was working on a high-speed switch code-named Vulcan. Yet another group in Yorktown was trying to solve both the problem of deploying these fast new RS/6000 workstations to the desktops of IBM workers and the system management headaches that come with LAN administration. This group developed a frame that could house 16 RS/6000 machines as well as management software called Agora to create a true turnkey LAN solution.

In December 1991, these independent efforts began to come together. HPSSDL was reorganized and renamed to HPSSL (the Development part of the name was dropped) under the leadership of Irving Wladawsky-Berger. HPSSL's mission was to ship a product in 12 months. Designing a system from scratch was out of the question given the time constraint; so, a task force was created that selected the necessary system components from available IBM technology. The RS/6000 Model 370 furnished the node technology. The Yorktown LAN consolidators provided their knowledge and experience in packaging the nodes. The Vulcan switch, now code-named Envoy (the origin of the `E` commands for the switch), was chosen over the ESCON interconnect, which was experiencing problems with excessive latency. Work from the ESCON interconnect experiment formed the basis for the first iteration of the Vulcan switch software. The total product was introduced to the marketplace as the SP1.

In September, 1993, Argonne National Laboratories in Argonne, Illinois, received shipment of the first SP1, a 128-node system. Cornell University in Ithaca, New York, bought a 512-node system shortly thereafter. Next came the petroleum industry. By the end of the year, 72 systems were installed around the world, all with scientific and technical customers.

Initially, IBM had no intention of positioning the SP1 in the commercial marketplace, but commercial customers started investigating IBM's SP. In the early 1990s, the death of the mainframe was accepted as conventional wisdom. Therefore, many large commercial enterprises were looking for alternatives to the mainframe to deploy new applications. IBM formed an application solutions group for the SP1, which, among other things, ported a parallel version of Oracle's database to the SP1. In 1994, SP development absorbed personnel from the discontinued AIX/ESA product, people who bolstered the system's manageability and helped spawn the Parallel System Support Programs (PSSP). By the end of 1994, the commercial segment accounted for 40 percent of all installed SPs. By the end of 1996, the share climbed to 70 percent.

The SP2 was announced in 1994. It incorporated new node types from Austin and a faster switch, code-named Trailblazer (the origin of the tb2 and tb3 nomenclature of the switch adapters). The SP2 had moved out from under the umbrella of the Large Systems Division and was its own enterprise within IBM. SP2 sales were strong: 352 systems were installed by the end of 1994, and 1,023 by the end of 1995.

In 1996, the SP2 was renamed to SP and formally became a product of the RS/6000 Division. It represents the high-end of the RS/6000 family. IBM secured a number of large SP contracts, of particular note the ASCI project of the U.S. Department of Energy. These contracts, coupled with the broad marketplace acceptance of the product, have fueled SP development. In 1996, IBM introduced a faster version of the Trailblazer switch, more than doubling the bandwidth of its predecessor, with new nodes, including Symmetric Multiprocessor (SMP) versions, and more robust and functional PSSP software.

7.2 SP Uses and Introduction

The RS/6000 SP high-performance and high-capacity system uses the power of parallel processing. Designed for performance and scalability, this system makes feasible the processing of applications characterized by large-scale data handling and compute intensity.

Customer uses include:

- Mission-critical commercial computing solutions such as ERP or e-commerce
- Business intelligence applications or data warehousing
- Server consolidation
- Collaborative computing such as Lotus Notes, Lotus Domino Server, Internet, intranet, extranet, and e-business
- Groupware application solutions.
- Scientific and technical computing users, including corporations, universities, and research laboratories

The SP database and computation scalability, critical for business intelligence applications including data warehousing, has led to many installations of greater than a terabyte of data.

Corporations, universities, and research laboratories use the SP system for leading-edge applications ranging from seismic processing, computational fluid dynamics, engineering analysis/design, and medical simulation.

7.3 Information Sources for the SP

The best place to start reading and learning about the SP are from these manuals:

- *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280
- *RS/6000 SP: Planning Volume 2, Control Workstation and Software Environment*, GA22-7281
- These can be downloaded for free from this URL:

http://www.rs6000.ibm.com/resource/aix_resource/sp_books

7.4 SP Hardware Components

The basic components of the RS/6000 SP are:

- Frames
- Internal Processor Nodes - High, Wide, and Thin
- Switches
- Control Workstation (CWS)

- SP Switch Router
- I/O Adapters
- SP Attached Servers
- Peripheral Devices

7.5 The Frame

The next section will outline the function of the frame, the components that are housed in the frame, and the history of the various frame models.

7.5.1 SP Frame Description

Strictly speaking, there are three types of frames:

- Tall frames
- Short frames
- Switch frames

The first two types are used to house processors or *nodes*, and they are usually simply called frames. The third one is used to house nodes and switches or Intermediate Switch Boards (ISB), which are described later in this chapter. This special type of frame can host up to eight switch boards.

SP nodes are mounted in either a tall or short frame. The frame spaces that the nodes fit into are called drawers. A tall frame has eight drawers, and a short frame has four drawers. Each drawer is further divided into two slots.

7.5.2 Contents of the Frame

The frame can contain the following devices:

- One slot can hold one thin node - maximum 16 nodes in a tall frame
- A wide node occupies one drawer (two slots) - maximum eight in a tall frame
- A high node occupies two drawers (four slots) - maximum of four in a tall frame
- An internal power supply is included with each frame
- Optional switch

7.5.3 SP Legacy Models and Differences of the New Frames

With the announcement in April 1998 of the new SP Model 500 (short frame) and SP Model 550 (tall frame), all previous models were withdrawn and are now referred to as *legacy* models. These legacy models are summarized in Table 76 and Table 77 on page 249.

Before this new frame offering was announced, the frame and the first node in the frame were tied together, forming a model. Each new node becoming available was potentially installable in the first slot of a frame, so a new model was born. With the new Model 500 and 550 offering, IBM simplified the SP frame options by decoupling the imbedded node from the frame offering. Therefore, when you order a frame, all you receive is a frame with the power supply units and a power cord. All nodes, switches, and other auxiliary equipment are ordered separately. This led to an increasing number of possible prepackaged configurations when more nodes became available.

In the SP Model 550, a new tall frame replaces the old one of the legacy models. The most noticeable difference between the new and old tall frame is the reduction in height. Another physical difference is the footprint.

All new designs are completely compatible with all valid SP configurations using older equipment. Also, all new nodes can be installed in any existing SP frame provided that the required power supply upgrades have been implemented in that frame. The reason for this is that the new nodes, such as the 375 MHz POWER3 Thin Node, introduced with PSSP 2.4, have a higher power consumption; hence there is a higher power requirement for the frame.

7.5.4 SP Legacy Models by Frame and Switch Types

Table 76 shows the SP legacy model categories determined by the type of switch installed, or not, and the height of the frame.

Table 76. *SP Legacy Models by Frame and Switch Types*

Model	Frame	Switch	Switch Frame	Nodes	Electrical Power
2Ax	Short	None	N/A	1 to 8	1 phase
3Ax	Short	SPS-8	No	1 to 8	1 phase
3Bx	Tall	SPS-8	No	1 to 8	3 phase
20x	Tall	None	N/A	1 to 64	3 phase
30x	Tall	SPS	No	1 to 80	3 phase
40x	Tall	SPS	Yes	65 to 512	3 phase

Model	Frame	Switch	Switch Frame	Nodes	Electrical Power
Notes: x - Designates node type included with frame Short - 124.5 cm (49 in) Tall - 200.7 cm (79 in) SPS - SP Switch SPS-8 - SP Switch 8-port model Switch Frame - Not available for < 64 nodes - Optional for 65-80 nodes - Required for > 81 nodes					

7.5.4.1 SP Legacy Models by Node Types

Table 77 provides a list of legacy models sorted by node type.

Table 77. SP Legacy Models by Node Types

Model	MHz	Form	Processor Type	I/O Bus	Available
xx1	62	Thin	Uni	MCA	No
xx2	66	Thin	Uni	MCA	No
xx3	66	Wide	Uni	MCA	No
xx4	66	Thin-2	Uni	MCA	No
xx5	77	Wide	Uni	MCA	No
xx6	112	High	SMP	MCA	No
xx7	135	Wide	Uni	MCA	No
xx8	120	Thin	Uni	MCA	No
xx9	200	High	SMP	MCA	No
xxA	160	Thin	Uni	MCA	No

7.5.5 SP Hardware Control and Supervision

Each frame (tall and short) has a supervisor card. This supervisor card connects to the Control Workstation using a serial link.

7.5.6 SP Frame Dimensions

Table 78 details the SP frame dimensions including the base.

Table 78. SP Frame Dimensions

Frame type	Height		Width		Depth		Weight	
	mm	in	mm	in	mm	in	kg	lbs
Tall Legacy models	2007	79	915	36	1118	44		
Tall 550 models	1925	75.8	922	36.3	1295	51	441 to 984	971 to 2165
Short 500 models	1245	49	711	28	1015	40	232 to 414	510 to 910

7.5.7 Frame Power Supplies

Tall frames come equipped with redundant (N+1) power supplies; if one power supply fails, another takes over. Redundant power is an option on short frames. These power supplies are self-regulating units. Power units with the N+1 feature are designed for concurrent maintenance; if a power unit fails, it can be removed and repaired without interrupting running processes on the nodes.

A tall frame has four power supplies. In a fully populated frame, the frame can operate with only three of these power supplies (N+1). Short frames come with two power supplies and a third optional one for N+1 support.

The power consumption depends on the number of nodes installed in the frame. For details, refer to *IBM RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280.

7.5.8 SP Frame Electrical Power Requirements

The electrical power specifications for the old and new power subsystems are summarized in Table 79.

Table 79. SP Frame Electrical Power Requirements

Power Requirements	Tall frame			Short frame	
	Legacy PDU	Legacy SEPBU	New Model 550	Legacy	New Model 500
SP Frame					
Power output (kW)	6.7	7	10.5	3.5	5.0
Utility loading (kVA) ¹	7.8	6.7	8.1	3.5	4.2
Phase	3	3	3	1	1
Current (Amps, low/high ² V)	24/16	24/16	35/20	24/-	30/-
Base Power Regulators ³	N/A	3	4	2	2
Computer room facility					
Conductor size (AWG, low/high ² V)	10/10	10/10	8/8	10	6
Circuit breaker (Amps, low/high ² V)	30/20	30/20	50/30	30/-	40/-
Connector (low voltage only ⁴)	460C9W ₅	430C9W ₅	460C9W ⁵	3933 ⁶	9C53U0 ₆
Receptacle (low voltage only ⁴)	460R9W ₅	430R9W ₅	460R9W ⁵	3753 ⁶	9R53U0 ₆
Notes: ¹ With thin nodes in every node position ² Low voltage: 200 - 240 V; High voltage: 380 - 415 V ³ With n+1 redundant frame power ⁴ High-voltage systems are wired to customer electrical service ⁵ Any IEC309 ⁶ Russell-Stoll					

Note

Visual inspection is needed to identify frames having an older Power Distribution Unit (PDU) versus the current Scalable Electrical Power Base Unit (SEPBU) so that the correct power upgrade can be ordered. For example, PDU frames have a circuit breaker and four 48 volt connectors, while SEPBU frames have no circuit breaker and eight 48 volt connectors. For detailed descriptions, refer to *IBM RS/6000 SP: Maintenance Information, Volume 1*, GC23-3903 and *IBM RS/6000 SP: Maintenance Information, Volume 2*, GC23-3904.

Power Cords

Also, the frame power upgrade requires a new power cord. Most frames have a 2 ¼ inch hole in the bottom of the frame, large enough to allow the power cord connector to go through. However, frames shipped in early October 1993 have a 2 inch hole the connector will not fit through. The power upgrade instructions states an electrician should be present to remove and reattach the external power cord connector so the new cord can be installed.

For more information about electrical power requirements for tall frames, refer to *IBM RS/6000 SP: Planning Volume 1, Hardware and Physical Environment GA22-7280*, Chapter 3 "Power and Electrical Requirements", tables titled:

- SEPBU Low-Voltage Requirements, 10.5 kW information
- SEPBU High-Voltage Requirements, 10.5 kW information
- AC Line Cord Specifications for SP Systems with the 10.5 kW SEPBU

For more information about electrical power requirements for short frames, refer to *IBM RS/6000 SP: Planning Volume 1, Hardware and Physical Environment, GA22-7280*, Appendix C "Model 2Ax and 3Ax Frame", tables titled:

- Model 2Ax and 3Ax Frame SEPBU, 10.5 kW information
- AC Line Cord Specifications for Model 2Ax or 3Ax with the 10.5 kW SEPBU

7.6 Internal Processor Nodes

The basic RS/6000 SP building block is the servers or processor *node*. Each node has a processor or processors, memory, disk drives, and PCI expansion slots for I/O and connectivity such as any other server. There are three types of internal SP nodes: thin, wide, and high. They may be mixed in a system and are housed in a frame. Figure 59 shows a mix of nodes in a frame.

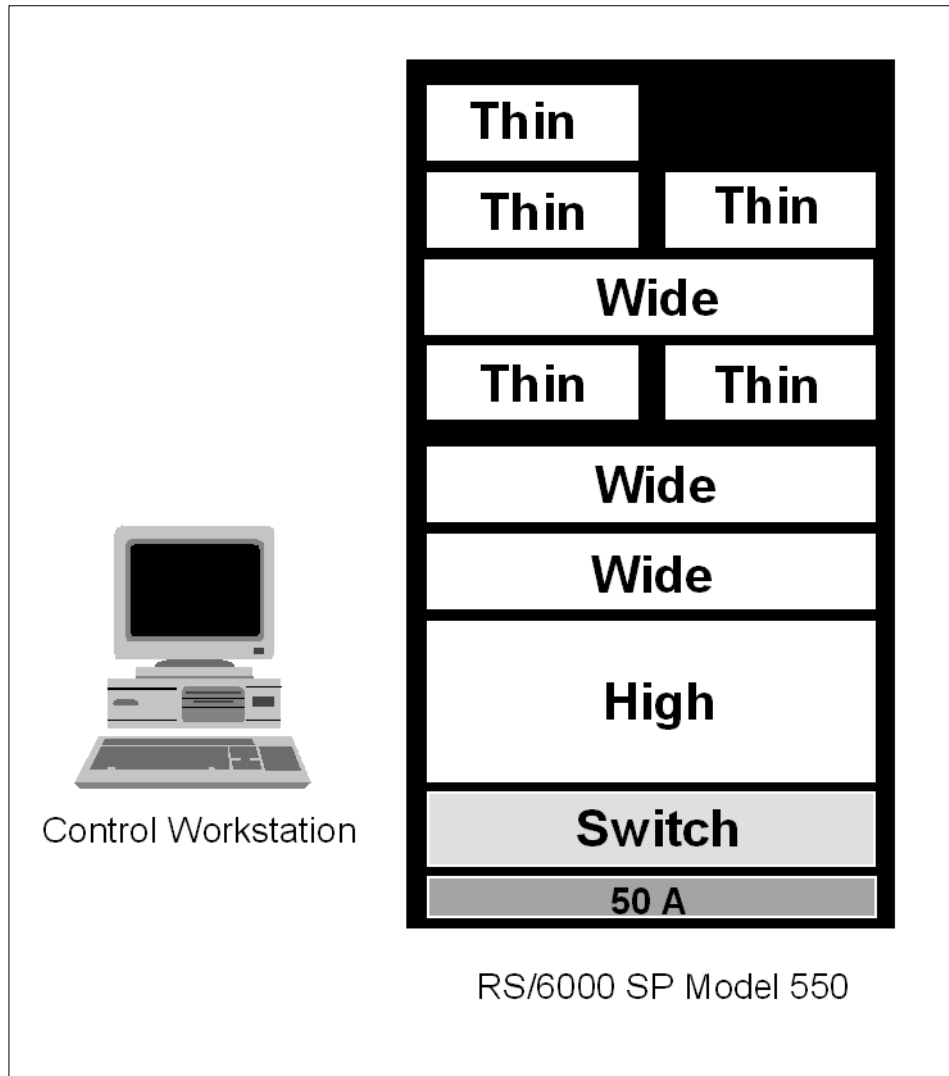


Figure 59. Internal SP Processor Nodes

7.6.1 Node Selection

At first glance, the number of different nodes and the different types of nodes can appear confusing, and to some it may not be obvious when to select a particular type of node within the SP system. This section looks at important considerations when selecting the best type of SP node for a particular purpose and environment.

7.6.1.1 Capacity

When selecting nodes for an SP system, first look at the required capacity for adapters, disk, and memory. This will help you decide whether a thin node has sufficient capacity. As you consider this area, take particular care to include all adapters that will be required—both now and in the near future.

Remember, for high availability solutions, you often need to have additional adapters for redundancy.

Remember also that the internal disk is used differently for different applications.

In the commercial world, it is usual to use the internal disk on each node only for the AIX operating system, paging space, application executables, and any temporary data. It is usual to store any critical customer data on external disks so that they can be physically connected to more than one AIX system to provide high availability in case of primary node failure. The internal disk cannot be accessed from another node if the node itself has failed.

This is in contrast to some scientific/technical applications where the volume of data is not as large, and high availability is not so important. In these cases, sometimes only using the internal disks is the appropriate solution.

Another factor to consider regarding capacity of nodes concerns memory. Due to the specific memory slots available (these are separate to PCI slots) in the SP nodes, not every combination of memory will be available—particularly if memory cards are already in use in an SP node.

The PCI nodes are available in three form factors: thin, wide, and high. A specific application example where the extra performance likely to be observed when using a PCI wide node as opposed to a PCI thin node would be when using the PCI node as a disk server node for a GPFS file system, where up to double the bandwidth could be achieved with a wide node rather than with a thin node.

7.6.2 SP Nodes at a Glance

Table 80 outlines the new POWER3-II 375 MHz Thin and Wide nodes that were announced in February 7, 2000.

Table 80. Standard Configurations for POWER3-II 375 MHz SMP Nodes

Node type	375 MHz POWER3 SMP Thin	375 MHz POWER3 SMP Wide
Feature code	2056	2057
Processor	375 MHz POWER3-II 2-way or 4-way	375 MHz POWER3-II 2-way or 4-way
Data/ Instruction (L1) cache	64 KB/32 KB	64 KB/32 KB
L2 cache/processor	8 MB	8 MB
Standard RAM	256 MB	256 MB
Maximum RAM	8 GB	8 GB
Disk/media bays	Two	Four
Storage	Internal/external boot disks Internal disks installed in pairs to support mirroring	Internal/external boot disks Internal disks installed in pairs to support mirroring
Max. internal storage	36.4 GB	109.2 GB
PCI bus speeds	33 MHz	33 MHz
PCI expansion slots	Two 32-bit	Eight 64-bit, two 32-bit
Integrated adapters	Integrated Ultra2 SCSI Ethernet (10/100 Mbps) High-speed serial port - 9pin D-shell Service processor	Integrated Ultra2 SCSI Ethernet (10/100 Mbps) High-speed serial port - 9pin D-shell Service processor
SP switch and adapter	MX2 Switch is supported (# 4023)	MX2 Switch is supported (# 4023)
ROLTP	44.0, 67.7	44.0, 80.0

Table 81 outlines the details of the POWER3 222 MHz High node that was announced September 13, 1999. See Section 7.6.3.1, “375 MHz POWER3 SMP High (# 2058)” on page 259 for more in depth information.

Table 81. POWER3 SMP 222 MHz High Node

Node type	375 MHz POWER3 SMP High
Feature code	2058
Processor	375 MHz POWER3-II 4-way, 8-way, 12-way, or 16-way
Data/instruction L1 cache	64 KB/32 KB
L2 cache/processor	8 MB
Standard RAM	1 GB
Maximum RAM	64 GB
Disk/media bays	2/26 with SP Expansion I/O Unit
AIX storage	Internal/external boot disks Internal disks installed in pairs to support mirroring
Maximum internal storage	946.4 GB with SP Expansion I/O Unit
PCI bus speed	33 MHz
PCI slots	5 PCI (one 32-bit and four 64-bit) and 8 PCI in the Expansion Unit (53-max)
Adapters	Integrated 10/100 Ethernet Integrated Ultra SCSI service processor
SP switch and adapter	MX2 Adapter
ROLTP	81.7, 163.7, 242.3, 319.3

Table 82 outlines the details of the 200 MHz Thin and Wide nodes that were announced February 1, 1999. Check with your IBM representative for their current availability.

Table 82. Standard Configuration for POWER3 200 MHz SMP Nodes

Node type	POWER3 SMP Thin	POWER3 SMP Wide
Feature code	2052	2053
Processor	200 MHz POWER3 1-way or 2-way	200 MHz POWER3 1-way or 2-way
L1 cache/processor	64 KB data, 32 KB instr.	64 KB data, 32 KB instr.
L2 cache/processor	4 MB	4 MB
Standard RAM	256 MB	256 MB
Maximum RAM	4 GB	4 GB
Memory bus width	128-bit	128-bit
Disk/media bays	Two	Four
AIX storage	Internal/external boot disks Internal disks installed in pairs to support mirroring	Internal/external boot disks Internal disks installed in pairs to support mirroring
Max. internal storage	36.4 GB	72.8 GB
Bus speeds		
For I/O adapters	132 MBps	132 and 264 MBps (triple bus)
For switch adapter	480 MBps	480 MBps
PCI expansion slots	Two 32-bit	Eight 64-bit, two 32-bit
Adapters	Integrated Ultra SCSI Ethernet (10/100 Mbps)	Integrated Ultra SCSI Ethernet (10/100 Mbps)
SP switch and adapter	122 MBps	122 MBps
ROLTP	10.5 (1-way), 21.0 (2-way)	10.5 (1-way), 21.0 (2-way)

Table 83 outlines the 332 MHz PowerPC 604e Thin and Wide nodes announced on May 20, 1998. These nodes use a 32-bit processor. Check with your IBM representative for their current availability.

Table 83. Standard Configuration for the 332 MHz PowerPC 604e Nodes

Node type	332 MHz SMP Thin	332 MHz SMP Wide
Feature code	2050	2051
Processor	332 MHz PowerPC 604e 2-way or 4-way	332 MHz PowerPC 604e 2-way or 4-way
L1 cache/processor	32 KB data, 32 KB instr.	32 KB data, 32 KB instr.
L2 cache/processor	256 KB	256 KB
Standard RAM	256 MB	256 MB
Maximum RAM	3 GB	3 GB
Memory bus width	128-bit	128-bit
Disk/media bays	Two	Four
AIX storage	Internal/external boot disks Internal disks installed in pairs to support mirroring	Internal/external boot disks Internal disks installed in pairs to support mirroring
Max. storage	36.4 GB	72.8 GB
Bus speeds		
For I/O adapters	132 MBps	132 and 264 MBps (triple bus)
For switch adapter	400 MBps	400 MBps
PCI expansion slots	Two 32-bit	Three 64-bit, seven 32-bit
Adapters	Integrated SCSI-2 F/W Ethernet (10 Mbps)	Integrated SCSI-2 F/W Ethernet (10 Mbps)
SP switch and adapter	122 MBps	122 MBps
ROLTP	17.9 (2-way), 32.8 (4-way)	17.9 (2-way), 32.8 (4-way)

7.6.3 High Nodes

Currently, there is only one high node available for the SP:

- 375 MHz POWER3 SMP High (# 2058)

If you need additional I/O slots or disk, you can order the following:

- SP Expansion I/O Unit

7.6.3.1 375 MHz POWER3 SMP High (# 2058)

The 375 MHz POWER3 SMP High node contains 375 MHz POWER3-II processors and is offered in 4-, 8-, 12-, and 16-way configurations. The new node, packaged in a high node configuration, can be added to an existing Model 550 frame, switchless or with a 16-port SP Switch. The SP Switch MX2 Adapter is supported on the new POWER3 SMP High Node.

The previously announced SP PCI Nodes can be converted to the new 375 MHz POWER3 SMP High Node. There is no upgrade path provided from RS/6000 SP Micro Channel nodes to the POWER3 SMP High Node.

POWER3-II Microprocessor

The POWER3-II microprocessor continues the POWER2 architecture tradition of bringing real solutions to IBM RS/6000 customers' high-performance computing needs, while adding 64-bit addressability, double-word integer operations, and SMP support. To satisfy compute-intensive requirements, the POWER3-II design contains a highly superscalar core, which is comprised of eight execution units capable of sustaining an execution rate of four instructions per cycle, a 32 KB instruction cache, 64 KB data cache, high bandwidth, independent interfaces to the 8 MB of L2 cache, and system memory.

Bus Description

Each processor card contains four POWER3-II microprocessors. Each processor employs 8 MB of cache on a private high speed bus. Four processors are attached to a set of node controller chips through independent point-to-point address and 16 byte wide. Each processor can issue up to one address transaction per cycle to the system.

The node controller chips themselves are attached to the rest of the system via a pair of unidirectional address buses and a pair of 16 byte wide bi-directional data buses. The node controller chips provide extensive buffering of addresses and data to improve system performance. Collectively, they also provide coherency management between all of the processors, memory, and I/O using a snoopy bus protocol.

Adapter Placement Restrictions

External disk booting has a prerequisite of an Ultra SCSI DE adapter (# 6207) or an SSA adapter (# 6225 or # 6230 if non-RAID) installed on the node.

Mandatory Prerequisites

Requires AIX Version 4.3.3.10 or later with IBM Parallel System Support Programs for AIX Version 3.2 or Version 3.1.1 with APAR IY11093.

In addition to the above, the new SMP High node requires APAR IY11426 with PSSP 3.2, or APAR IY11425 with PSSP V3.1.1 if there are 64 GB of memory with SP Switch MX2 Adapter installed.

Standard Features

There are four CPU slots, four memory card slots, two storage bays, five PCI slots, and six I/O connections

It includes integrated 10/100 Mbps Ethernet, an Ultra SCSI interface, and supports up to 16 GB of memory and 72.8 GB of internal disk.

Memory and L2 cache are protected with single bit correct, double bit detect ECC. L1 cache is parity protected. Memory incorporates scrubbing and supports continued operation with a full memory chip failure as further protection.

Both the POWER3 SMP High Nodes and the SP Expansion I/O Unit incorporate N+1 power and N+1 cooling. The Expansion Units also provide redundant links to the High Node.

Disk mirroring is standard on both the High Node and the Expansion Unit. It provides redundant storage allowing continued operation in the presence of a disk failure. Disks are shipped in pairs. If external boot feature code (# 9121 SCSI or # 9124 SSA) is selected, zero internal disks are allowed

There is one CPU card, a 4-way and four available CPU card slots.

CPU and Memory RepeatGuard capability is provided. This function checks for excessive soft or hard fails at boot time and deconfigures a faulty memory bank or processor for deferred repairs. The soft fail threshold is user selectable.

Performance

The 375 MHz POWER3 SMP High Node's superior performance characteristics are the result of clearly articulated design objectives gleaned from in-depth knowledge of challenging customer applications and years of experience in designing the RS/6000 family of POWER and POWER2

processors. Consistent with its POWER and POWER2 heritage, the POWER3 SMP High Node distinguishes itself from its competitors in its ability to deliver main memory bandwidth, a crucial attribute to many applications on the computing frontier. These applications can achieve a quantum leap in performance on the High Node and, in turn, provide new scientific insights and competitive advantages to their users.

The POWER3 SMP High Node excels in data delivery; indeed, this was its most significant and most challenging design objective. The most sophisticated features of its design are all integral pieces of its data delivery system: dual load/store execution units, an interleaved data cache, multiple-outstanding cache-miss support, wide data paths to memory and its L2 cache, hardware data prefetch, node controllers, and a non-blocking data cross point switch.

Complimenting the POWER3 SMP High Node's excellent memory bandwidth are floating-point and fixed-point execution units that can sustain remarkable rates of computation.

Memory Requirements and Options

Memory is packaged on up to four cards connected to the four memory ports in the data cross point switch. It is recommended to install four memory cards with 2 GB of DIMMS one each to maximize bandwidth. Memory cards utilize 128, 256, or 512 MB SDRAM modules. Each card can contain up to 16 GB of memory or up to 64 GB for the complete node. One GB of memory comes in the base system design. Memory can be added in 1 GB increments with the addition of cards and pluggable industry standard DIMMs. This provides a highly configurable and upgradable offering, which will grow with the application requirements.

The option to order three base memory cards is *not* supported.

To maximize the bandwidth of the memory, it is highly recommended to install 2 GB of 8 X 128 MB DIMMs (2 X # 4133) on each of the four base memory cards.

Table 84 outlines the various memory options available to the 375 MHz POWER3 High node.

Table 84. POWER3 SMP 222 MHz High Node Memory Options

Quantity	Feature Code	Description
4 - default amount	4880	Memory Expansion Feature
8 - default amount	4133	1 GB (8x128 MB) DIMMs

Quantity	Feature Code	Description
0 - 16	4134	2 GB (8x256 MB) DIMMs
0 - 8	4402	2 GB (16x128 MB) DIMMs
0 - 4	4403	3 GB (24x128 MB) DIMMs
0 - 4	4404	4 GB (32x128 MB) DIMMs
0 - 8	4412	2 GB (16x256 MB) DIMMs
0 - 4	4413	3 GB (24x256 MB) DIMMs
0 - 4 m	4414	4 GB (32x256 MB) DIMMs
0 - 16	4421	4 GB (8x512 MB) DIMMs
0 - 8	4422	8 GB (16x512 MB) DIMMs
0 - 4	4423	12 GB (24x512 MB) DIMMs
0 - 4	4424	16 GB (32x512 MB) DIMMs

Disk Requirements and Options

Table 85 outlines the options of disk offered internally to the POWER3 High node and potentially in the Expansion I/O drawer.

Table 85. Internal Disk on POWER3 222 MHz High Node or its Expansion I/O

Feature Code	Description	POWER3 High Node (Min. and Max. Pair disk)	Expansion I/O Unit (Min. and Max. pair disk)
2909	9.1 GB Ultra SCSI (pair)	1-1	
2918	18.2 GB Ultra SCSI (pair)	0-1	
3804	9.1 GB 10K RPM Ultra SCSI (pair)	0-1	0-2
3810	18.2 GB 10K RPM Ultra SCSI (pair)	0-1	0-2
3820	36.4 GB 10K RPM Ultra SCSI (pair)	0-1	0-2

Switch Adapter Requirements

Supports SP Switch MX2 Adapter (#4023) with 16-port SP Switch only.

7.6.3.2 Expansion I/O Drawer (# 1550)

The SP Expansion I/O Unit provides an extension to the POWER3 SMP High Node with eight PCI hot-pluggable 64-bit PCI slots running at 33 MHz and four hot swappable Ultra SCSI internal disk storage bays. A 9.1 GB mirrored disk pair is supported for SCSI and SSA adapters, and an 18.2 GB or 36.4 GB mirrored disk pair is supported for SCSI adapters only. The SP Expansion I/O Unit may be installed in a frame with other SP nodes or in an SP frame that contains only the SP Expansion I/O Unit.

- Supports only the POWER3 SMP High node.
- Cannot co-exist with a PCI node in the same SP node drawer.
- Requires an SCSI adapter or an SSA adapter for internal disk storage.
- SCSI and SSA disk storage cannot be mixed within the same SP Expansion I/O Unit.
- Up to 6 of these Units can be attached to a POWER3 SMP High node.
- A maximum of 16 of these can be attached to one frame.
- A SCSI disk pair requires an Ultra SCSI SE adapter (# 6206). An SSA disk pair requires an Advanced SerialRAID adapter (# 6225 or 6230).
- An air baffle (# 9612) is provided for each empty disk pair.
- Equal in size to one thin node.
- The Expansion Unit provides *hot-swappable* capability for both SCSI and SSA disks, which enables maintenance concurrent with node operation. This reduces down time for maintenance.

In addition, there are six remote I/O (RIO) connections on the POWER3 High node. These connections allow SP Expansion I/O Expansion Units to be connected to the High Node providing incremental growth for applications requiring more I/O connectivity. The RIO ports operate full duplex at 250 MBps each direction with cables up to 15 meters long. This provides outstanding I/O bandwidth to each RIO expansion node. With six SP Expansion I/O Units connected, a user can obtain a node complex with 53 PCI slots and 26 DASD bays. This combination of High Node and SP Expansion I/O Units into a node complex provides the highest I/O configuration yet on the RS/6000 SP system.

PCI Hot Pluggable Adapters in SP Expansion I/O Units

Hot Plug capabilities are now supported on the RS/6000 SP Expansion I/O Units for PCI adapters listed below. This new function enables the maintenance of these adapters without powering off the SP Expansion I/O Unit or the associated SP POWER3 SMP High Node.

This new function requires AIX Version 4.3.3 with APAR IY06844 and will be included on all SP Expansion I/O Units after February 25, 2000. Customers who have an SP Expansion I/O Unit shipped prior to this date can easily add this function by updating the firmware and installing AIX Version 4.3.3 with ARAP IY06844. The required firmware is available through the IBM Web site at:

<http://www.austin.ibm.com/support/micro/download.html>

The supported PCI adapters are as listed in Table 86.

Table 86. Hot Plug PCI Adapters for 375 MHz POWER3-II Wide Node

Feature Code	Description
2732	Short Wave Serial HIPPI
2733	Long Wave Serial HIPPI
2741	FDDI SK-NET LP SAS
2742	FDDI SK-NET LP DAS
2743	FDDI SK-NET UP SAS
2751	S/390 ESCON Channel
2920	Token ring Auto Lanstream
2943	WAN RS232/RS422 8 Port
2944	WAN RS232/RS422 128 Port
2947	ARTIC960HX 4-Port Multiprotocol
2962	2-Port Multiprotocol
2963	TURBOWAYS 155 UTP ATM
2968	Ethernet 10/100 MB
2969	Gigabit Ethernet
2985	Ethernet BNC/RJ-45
2987	Ethernet AUI/RJ-45
2988	TURBOWAY ATM 155 MMF
4959	Token ring
6204	Universal Differential Ultra SCSI
6205	Dual Channel Ultra2 SCSI

Feature Code	Description
6206	Ultra SCSI SE
6207	Ultra SCSI DE
6225	Advanced SerialRAID Adapter
6227	Gigabit Fibre Channel for PCI Bus
6230	Advanced SerialRAID Plus Adapter

7.6.4 Wide Nodes

Currently there are three different types of Wide nodes available:

- 375 MHz POWER3 SMP Wide (# 2057)
- 332 MHz SMP Wide (# 2051)

The following sections outline the individual nodes and their abilities and potential components.

7.6.4.1 375 MHz POWER3 SMP Wide (# 2057)

The new 64-bit, 375 MHz POWER3 SMP Wide node offers excellent integer and floating-point performance and price/performance. These nodes are available in 2- and 4-way POWER3-II CPU configurations. When fully configured with four POWER3-II CPUs, these nodes offer up to two times more integer performance than 332 MHz SMP Thin or Wide Nodes. The strong integer performance makes these nodes ideal for business-critical applications, such as data mining and warehousing, Enterprise Resource Planning(ERP), collaborative computing, and Web serving.

These offerings also expand the range of node choices for customers, reducing administrative overhead and increasing productivity by consolidating UNIX servers. In addition to handling existing 32-bit commercial applications, these nodes also provide support for applications that benefit from 64-bit technology, such as data warehousing, multimedia Web serving, and ERP. The new fully configured 4-way nodes also have up to three times more floating-point performance than POWER3 SMP Thin or Wide Nodes. This powerful floating-point performance makes these nodes ideal for technical computing applications, such as seismic processing, engineering, analysis, scientific research, chemical modeling, and electronic design automation.

These new nodes can be added to existing Model 5xx frames. The current 332 MHz SMP and POWER3 SMP Thin and Wide Nodes can be upgraded to the new 375 MHz POWER3 SMP Nodes. These upgrade paths help preserve

your investment in SP PCI (Peripheral Component Interconnect) nodes. Users of older MCA (Micro-Channel Architecture) nodes cannot upgrade to the 375 MHz POWER3 SMP Nodes.

This node utilizes the 64-bit 375 MHz POWER3-II microprocessor with 8 MB of Level 2 (L2) cache per processor. The 375 MHz POWER3 SMP Nodes, packaged in Thin and Wide configurations, can be installed in an SP system, having appropriate electrical power, for a maximum of 128 nodes. Systems greater than 128 nodes are available via special bid.

This node can be functionally compared to the Model 270.

Note

Applications utilizing a high percentage of node resources, such as memory, processor cycles, and I/O bandwidth, may not be able to utilize as many PCI slots, and still obtain high throughput, as applications using a low percentage of node resources.

Bus Description

The system bus is controlled by a highly specialized set of custom chips. One handles addressing and synchronization, the other moves data to and from the processor (the 6XX bus), memory (using memory bus), and the I/O (I/O bus).

The 6XX bus is a 128-bit bus. It is optimized for high-performance and multiprocessing performance. The bus is fully parity checked and each memory request is range checked and positively acknowledged for error detection. Any error will cause a machine check condition and is logged in the AIX error log.

The 6XX and memory buses operate at the same speed, 128-bit width, and have the same throughput. Their speed is automatically determined by the speed of the processor installed. Data and address buses operate independently in true split transaction modes and are pipelined so that the new requests may be issued before previous requests are completed.

The POWER3-II 375 MHz wide node is compliant with Revision 2.1 of the PCI specifications and implements dual PCI bridge chips in a per configuration. One PCI bridge chip provides a 32-bit interface operating at 33 MHz for two PCI slots and the other PCI bridge chip implements a 64-bit bus operating at 50 MHz for eight PCI slots.

Adapter Placement Restrictions

The maximum of # 6225 and # 6230 combinations is two per bus and total of six per 375 MHz POWER3 SMP Wide Node (# 2057).

The 36.4 GB disk drive (# 3820) is only supported on 375 MHz POWER3 SMP Wide Node (# 2057).

Consult the PCI Adapter Placement Guide for individual details.

Mandatory Prerequisites

Model 2xx, 3xx, and 4xx frames and related expansion frames must have the electrical power subsystem upgrade (# 9932, 9933, or 9934) before the first PCI node can be added to a currently installed frame.

Limitations:

- SP Switch MX Adapter (# 4022) is not supported.
- Micro Channel (MCA) adapters are not supported.
- The High Performance Switch (# 4010 or # 4007) is not supported.

The following must be installed on the POWER3 SMP Node and the SP Control Workstations:

- AIX Version 4.3.3 with APAR IY06844 or later
- IBM Parallel System Support Programs for AIX (PSSP) Version 3.1.1 or later.

Standard Features

The summary of standard features is as follows:

- 2- or 4-way SMP design (within two CPU card slots)
- 375 MHz POWER3-II 64-bit processors
- 8 MB of L2 cache per processor
- 256 MB of ECC SDRAM memory expandable to 8 GB
- Four disk drive bays per wide node
- Pair of 9.1, 18.2, or 36.4 Ultra SCSI or 10K RPM disks
- Ten PCI slots (two 32-bit, eight 64-bit) per wide node
- Integrated 10/100 Ethernet
- Integrated Ultra2 SCSI interface (supports Ultra2 SCSI internal disks)
- Integrated high-speed serial port (S2, a 9-pin D-shell)
- Integrated service processor

Processor Requirements and Options

A 2-way 375 MHz processor (# 4444) is available on this node. A maximum of two of these processor boards can be ordered. Included with each processor is 8 MB of L2 cache.

CPU Gard, or processor run-time deconfiguration, is a RAS enhancement that can remove error-prone CPUs from the system configuration during run-time and at subsequent node boots. This minimizes node failures and data integrity exposures due to a processor that is experiencing repeated run-time recoverable internal errors.

CPU Gard is implemented in the system firmware, service processor firmware, and AIX operating system (4.3.3 + APAR IY06844 Version or later). When a processor's internal recoverable errors reach a predefined threshold, the firmware notifies the AIX operating system which migrates all software processes and interrupts to another processor and puts the faulty processor in a *stop state*. The processor remains offline for subsequent reboots via the CPU Boot Time Deconfiguration function, until the faulty hardware is replaced. The user can enable and disable CPU Gard via the AIX system management services.

Memory Requirements and Options

The 375 MHz POWER3 SMP Nodes, offered in 2- or 4-way configurations, support up to 8 GB of memory (two times that of the current POWER3 SMP Thin and Wide Nodes), a range of PCI adapters, and include integrated 10/100 Mbps Ethernet and Ultra2 SCSI interfaces.

Table 87 on page 268 outlines the potential memory DIMMS available.

Table 87. *Minimum and Maximum Memory*

Feature Code	Description	Minimum	Maximum
4098 - default	Memory Expansion Feature	2	2
4110 - default	256 MB (2 x 128 MB) SDRAM DIMM	1	16
4119	512 MB (2 X 256 MB DIMM)	0	16
4133	1 GB Memory (8 X 128 MB DIMM)	0	4
4402	2 GB Memory (16 X 128 MB DIMM)	0	2

Feature Code	Description	Minimum	Maximum
4403	3 GB Memory (24 X 128 MB DIMM)	0	1
4404	4 GB Memory (32 X 128 MB DIMM)	0	1
4412	4 GB Memory (16 X 256 MB DIMM)	0	2
4413	6 GB Memory (24 X 256 MB DIMM)	0	1
4414	8 GB Memory (32 X 256 MB DIMM)	0	1

Disk Requirements and Options

The 375 MHz POWER3 SMP Wide Node can support up to 109.2 GB of mirrored internal disk. The option of ordering no internal disks is available for customers who choose to use external disk storage and is done by not selecting internal storage and selecting (# 9121) for external SSA boot or (# 9124) for external SCSI Boot. Configuring these nodes for mirroring is strongly recommended.

Table 88 on page 269 outlines the potential disks, which only come in pairs, that can be installed on the 375 MHz POWER3 Wide Node.

Table 88. Potential Disk pairs in 375 MHz Wide Node

Feature Code	Descriptions
2909	9.1 GB Ultra SCSI DISK Drive Pair
2918	18.2 GB Ultra SCSI Disk Drive Pair
3804	9.1 GB 10K RPM Ultra SCSI Disk Pair
3810	18.2 GB 0K RPM Ultra SCSI Disk Pair
3820	36.4 GB 10K RPM Ultra SCSI Disk Pair

Switch Adapter Requirements

POWER3 375 MHz SMP Nodes are connected to the SP Switch using the SP Switch MX2 Adapter (#4023), which attaches directly to the processor mezzanine (MX) bus, providing increased bandwidth for TCP/IP, UCP/IP, and user-space communications between nodes. It does not require a PCI slot, which leaves one more adapter slot available for application use.

In a system configured with an SP-8 switch, any MES that increases the form factor of a node, such as thin to wide node or wide to high node, will require *all* nodes that are above the node being upgraded to be deleted via the `spdelnode` command and reinstalled. Additionally, the SPS switch cables on *all* the nodes above the one being upgraded need to be recabled to represent the new switch node positions.

Dynamic Processor Deallocation

Another high availability feature is processor deallocation. This feature allows the system to completely isolate a processor that has been determined to be operating below an acceptable level of reliability. To accomplish this, the system monitors processor status in real time to determine when the number of recoverable errors has exceeded a specified threshold. When the threshold is exceeded, the system sends a report to the operating environment that a processor is unreliable. The AIX operating environment then re-routes all interrupts and processes to other processors and disables the unreliable processor complex without interrupting end user services. Firmware level 991216 must be installed on the 375 MHz Wide Node. Since it is a feature that was introduced to AIX in February 7, 2000, the AIX section of this book discusses this function further.

The Web site where microcode/firmware can be downloaded from is located at:

<http://www.austin.ibm.com/support/micro/download.html>

7.6.4.2 332 MHz SMP Wide (# 2051)

332 MHz SMP Wide nodes (# 2051) have PCI bus architecture and use either two or four 332 MHz PowerPC processors per node. These nodes are functionally equivalent to an IBM RS/6000 7025-F50 workstation. The SP system must be operating at PSSP 2.4 (or later) to use these nodes.

The 332 MHz SMP Wide node occupies one full drawer; therefore, eight SMP wide nodes can be housed in a tall frame. SMP wide nodes can be placed in the first node slot of a frame without requiring additional nodes. However, uniprocessor wide nodes in the first node slot still require an additional filled node drawer in that frame.

For electromagnetic compliance, these nodes are housed in an SMP Enclosure.

This SMP enclosure (# 9930) is automatically included when you order a 332 MHz SMP wide node.

If you are going to mount a 332 MHz SMP wide node into an older 2.01 m or 1.25 m frame, a power system upgrade is necessary. However, once you have done the power system upgrade, these nodes are fully compatible with all existing SP hardware except for High Performance Switches.

332 MHz SMP wide nodes are not compatible with High Performance Switches (# 4010 and # 4007).

Bus Description

The 332 MHz SMP Wide node PCI bus is divided into three logical groups of PCI slots. The first slot group (slots I2 and I3) is composed of the two 32-bit slots residing on the CPU side of the 332 MHz SMP wide node, and the second and third group reside on the I/O side of the node. Both the second and third group have four PCI slots each. The second group (slots I1 through I4) has three 64-bit slots and a single 32-bit slot. The third group (slots I5 through I8) is made up of the last four 32-bit slots on the I/O side of the node. The third group is a physical extension on the second group. The I1 slot on the CPU side of the node is reserved for the optional SP Switch MX2 Adapter.

Previously installed 332 MHz SMP Wide nodes may have a withdrawn SP Switch MX Adapter in the CPU side I1 slot.

Adapter Placement Restrictions

With few exceptions, the 10 PCI slots in the 332 MHz SMP Wide node can be used for any valid RS/6000 SP PCI system adapter. While most PCI adapters will function in any 332 MHz SMP Wide node slot, the following adapters cannot be placed in any one of the third group of PCI slots:

- S/390 ESCON (# 2751)
- ARTIC960Hx 4-Port selectable (# 2947)
- 2-port Multiprotocol (# 2962)
- ATM 155 UTP (# 2963)
- Gigabit Ethernet - SX (# 2969)
- ATM 155 MMF (# 2988)
- Ultra SCSI SE (# 6206)
- Ultra SCSI DE (# 6207)
- SSA RAID5 (# 6215)
- ARTIC960RxD Quad Digital Trunk (# 6310)

To achieve the best performance with SSA RAID and Ultra SCSI DASD subsystems, the following adapters for these devices should be distributed evenly across the two recommended PCI slot groups:

- SSA RAID5 (# 6215)
- Ultra SCSI SE (# 6206)
- Ultra SCSI DE (# 6207)

To avoid performance degradation, the following adapters should not be placed in slots I5, I6, I7, or I8 in 332 MHz SMP wide nodes:

- FDDI SK-NET LP SAS (# 2741)
- FDDI SK-NET LP DAS (# 2742)
- FDDI SK-NET UP SAS (# 2743)
- 10/100 MB Ethernet (# 2968)
- SCSI-2 F/W single-ended (# 6208)
- SCSI-2 F/W differential (# 6209)

For similar reasons, if two S/390 ESCON adapters (# 2751) are used in this node, one adapter must be placed in the CPU bus, and the other adapter must be placed in the first I/O bus.

Mandatory Prerequisites

332 MHz SMP Wide nodes occupy one full node drawer. These nodes are symmetrically configured for memory, DASD, and adapters. Up to eight 332 MHz SMP wide nodes may be installed in one tall frame and up to four in a short frame. The mandatory prerequisites are:

- PSSP 2.4 (or later) on the Control Workstation, backup nodes, and processor node
- Two processors (mounted in one slot)
- 256 MB of memory
- 9.1 GB of DASD (with internal booting)
- An upgraded power system on older frames

Standard Features

Each 332 MHz SMP Wide node is functionally equivalent to an RS/6000 7025-F50 and has:

- Two processor slots allowing a maximum of four processors per node
- Two memory slots supporting up to 3 GB of memory

- Ten PCI slots for communication adapters (seven 32-bit and three 64-bit)
- A dedicated mezzanine (MX) bus slot for an optional switch adapter
- Integrated 10BaseT/10Base2 Ethernet (only one port may be used at one time)
- Four DASD bays supporting up to 72.8 GB of disk storage
- Integrated SCSI-2 Fast/Wide
- Standard Service Processor
- External nine-pin RS-232 on the planar S2 port

This connection has active heartbeat and is available for customer applications.

Processor Requirements and Options

SMP wide nodes require a minimum of two 332 MHz PowerPC processors mounted on one card. However, you can order an additional processor card (# 4320) to configure the node with a total of four CPUs. Table 89 provides the available processor options for the SMP nodes.

Table 89. Processor Options for 332 MHz SMP Wide Nodes

Feature Code	Multiplier	Description	Comments
4320	X 1	One processor card with two CPUs	Minimum required
4320	X 2	Two processor cards with two CPUs each (four CPUs total)	Maximum allowed

Memory Requirements and Options

332 MHz SMP Wide nodes have two memory cards and require a minimum of 256 MB of memory. These nodes support a maximum of 3 GB of memory. Memory is supplied by 128 MB DIMMs that must be mounted in pairs (256 MB increments).

The memory cards are not required to be configured symmetrically. Each card has the capacity to mount 2 GB of DIMMs; however, only 3 GB are addressable per node. Memory cards and DIMMs are not interchangeable

between SMP and non-SMP wide nodes. Table 90 provides the available memory features for the SMP nodes.

Table 90. Memory Features for 332 MHz SMP Wide Nodes

Feature Code	Description	Minimum Node Requirement	Maximum Allowed Per Node
4093	Base Memory Card	2	2
4110	One Pair of 128 MB DIMMs (256 MB total)	1 pair	6 pairs
4133	1 GB (8 x 128 MB) DIMMs	0	12
4402	2 GB (16 x 128 MB) DIMMs	0	1
4403	3 GB (16 x 128 MB) DIMMs	0	1

Disk Requirements and Options

332 MHz SMP Wide nodes can have up to four internal disks attached through an integrated SCSI-2 network. The 332 MHz SMP Wide node can have either no internal DASD (with external booting) or from 4.5 GB to a maximum of 72.8 GB of internal disk storage. External storage devices can be accessed through an optional Ultra SCSI Adapter (# 6207) or SCSI-2 Adapter (# 6209).

Optional direct access storage devices are available as follows:

- 9.1 GB Ultra SCSI disk drive pair (# 2909)
- 18.2 GB Ultra SCSI disk drive pair (# 2918)

This node does not require special cables or adapters to mount internal disks. However, the 332 MHz SMP wide node has an option (# 1241) that provides an independent SCSI hookup. It accomplishes the following:

- Eliminates the DASD controller as a single point of failure during mirroring
- Increases disk performance
- Balances disk loading

The (# 1241) option requires either an (# 6206) SCSI-2 Ultra/Wide Adapter PCI or (# 6208) SCSI-2 Fast/Wide Adapter 4-A PCI as a PCI-type SCSI adapter.

Switch Adapter Requirements

The switch adapter for SMP wide nodes does not occupy a PCI slot. Instead, the switch adapter for these nodes is installed into the Mezzanine (MX) bus. The MX bus connects the I/O planar with the system planar. Placing the

switch adapter in the MX bus enables switch traffic to proceed at higher bandwidths and lower latencies.

In switch-configured systems, 332 MHz SMP wide nodes require the following switch adapter:

- SP Switch MX2 Adapter (# 4023)

332 MHz SMP Wide Node Switch Restrictions

The 332 MHz SMP Wide node is not compatible with the older High Performance series of switches.

If an SMP wide node is going to be placed into an SP system configured with a switch, that switch must be either an SP Switch or an SP Switch-8.

Switch adapters for SMP wide nodes are not interchangeable with switch adapters used on uniprocessor wide nodes.

7.6.5 Thin Nodes

This section describes the various offerings of the thin nodes currently available. They are:

- 375 MHz POWER3 SMP Thin Node (# 2056)
- 332 MHz SMP Thin Node (# 2050)

7.6.5.1 375 MHz POWER3-II Thin (# 2056)

The new 64-bit, 375 MHz POWER3 SMP Thin node offers excellent integer and floating-point performance and price/performance. These nodes are available in 2- and 4-way POWER3-II CPU configurations. When fully configured with four POWER3-II CPUs, these nodes offer up to two times more integer performance than 332 MHz SMP Thin or Wide Nodes. The strong integer performance makes these nodes ideal for business-critical applications, such as data mining and warehousing, Enterprise Resource Planning (ERP), collaborative computing, and Web serving.

These offerings also expand the range of node choices for customers reducing administrative overhead and increasing productivity by consolidating UNIX servers. In addition to handling existing 32-bit commercial applications, these nodes also provide support for applications that benefit from 64-bit technology, such as data warehousing, multimedia Web serving, and ERP. The new fully configured 4-way nodes also have up to three times more floating-point performance than POWER3 SMP Thin or Wide Nodes. This powerful floating-point performance makes these nodes ideal for technical

computing applications, such as seismic processing, engineering analysis, scientific research, chemical modeling, and electronic design automation.

These new nodes can be added to existing Model 5xx frames. The current 332 MHz SMP and POWER3 SMP Thin and Wide Nodes can be upgraded to the new 375 MHz POWER3 SMP Nodes. These upgrade paths help preserve your investment in SP PCI (Peripheral Component Interconnect) nodes. Users of older MCA (Micro-Channel Architecture) nodes cannot upgrade to the 375 MHz POWER3 SMP Nodes.

This node utilizes the 64-bit 375 MHz POWER3-II microprocessor with 8 MB of Level 2 (L2) cache per processor. The 375 MHz POWER3 SMP Nodes, packaged in Thin and Wide configurations, can be installed in an SP system, having appropriate electrical power, for a maximum of 128 nodes. Systems greater than 128 nodes are available via special bid.

This node can be functionally compared to the Model 270.

Note

Applications utilizing a high percentage of node resources, such as memory, processor cycles, and I/O bandwidth, may not be able to utilize as many PCI slots, and still obtain high throughput, as applications using a low percentage of node resources.

Bus Description

The system bus is controlled by a highly specialized set of custom chips. One handles addressing and synchronization, the other moves data to and from the processor (the 6XX bus), memory (using memory bus), and the I/O (I/O bus).

The 6XX bus is a 128-bit bus. It is optimized for high-performance and multiprocessing performance. The bus is fully parity checked and each memory request is range checked and positively acknowledged for error detection. Any error will cause a machine check condition and is logged in the AIX error log.

The 6XX and memory buses operate at the same speed, 128-bit width, and have the same throughput. Their speed is automatically determined by the speed of the processor installed. Data and address buses operate independently in true split transaction modes and are pipelined so that the new requests may be issued before previous requests are completed.

The POWER3-II 375 MHz Thin node is compliant with Revision 2.1 of the PCI specifications and implements dual PCI bridge chips in a configuration. One PCI bridge chip provides a 32-bit interface operating at 33 MHz for two PCI slots.

Each POWER3 SMP Node drawer can contain one wide node, one thin node, or two thin nodes. When one POWER3 or 375 MHz POWER3-II SMP Thin Node is installed in a drawer, a node enclosure (#9930) and a filler plate (#9931) are required. A POWER3 Thin Node and a 375 MHz POWER3-II Thin Node can share the same node drawer. All previously announced SP PCI Thin and Wide nodes can be converted to a single 375 MHz SMP thin or wide node.

Adapter Placement Restrictions

Consult the PCI Adapter Placement Guide for individual details.

All SP PCI adapters are available on the 375 MHz POWER3 SMP Nodes with the same configuration rules and limitations as the corresponding POWER3 SMP Thin Node and Wide Node. # 2920, 6207, 6209, 6215, and 6225 are supported on the 375 MHz POWER3 SMP for MES upgrade orders when the feature is moved from an older node to a new 375 MHz POWER3 node, but not for orders of new 375 MHz POWER3 SMP nodes.

Mandatory Prerequisites

Model 2xx, 3xx, and 4xx frames and related expansion frames, must have the electrical power subsystem upgrade (# 9932, 9933, or 9934) before the first PCI node can be added to a currently installed frame.

- SP Switch MX Adapter (# 4022) is not supported.
- Micro Channel (MCA) adapters are not supported.
- The High Performance Switch (# 4010 or # 4007) is not supported.

The following must be installed on the POWER3-II SMP Node and the SP Control Workstations:

- AIX Version 4.3.3 with APAR IY06844 or later
- IBM Parallel System Support Programs for AIX (PSSP) Version 3.1.1 or later.

Standard Features

The summary of standard features is as follows:

- 2- or 4-way SMP design (within two CPU card slots)
- 375 MHz POWER3-II 64 bit processors

- 8 MB of L2 cache per processor
- 256 MB of ECC SDRAM memory expandable to 8 GB
- Four disk drive bays per thin node
- Pair of 9.1, 18.2, or 36.4 Ultra SCSI or 10K RPM disks
- Two 32-bit PCI slots per thin node
- Integrated 10/100 Ethernet
- Integrated Ultra2 SCSI interface (supports Ultra2 SCSI internal disks)
- Integrated high-speed serial port (S2, a 9-pin D-shell)
- Integrated service processor

Processor Requirements and Options

A 2-way 375 MHz processor (# 4444) is available on this node. A maximum of two of these adapters can be ordered. Included with each processor is 8 MB of L2 Cache.

CPU Gard, or processor run-time deconfiguration, is a RAS enhancement which can remove error-prone CPUs from the system configuration during run-time and at subsequent node boots. This minimizes node failures and data integrity exposures due to a processor that is experiencing repeated run-time recoverable internal errors.

CPU Gard is implemented in the system firmware, service processor firmware, and AIX operating system (4.3.3 + APAR IY06844 version or later). When a processor's internal recoverable errors reach a predefined threshold, the firmware notifies the AIX operating system which then migrates all software processes and interrupts to another processor and puts the faulty processor in *stop state*. The processor remains offline for subsequent reboots via the CPU Boot Time Deconfiguration function, until the faulty hardware is replaced. The user can enable and disable CPU Gard via the AIX system management services.

Memory Requirements and Options

The 375 MHz POWER3-II SMP Nodes, offered in 2- or 4-way configurations, support up to 8 GB of memory (two times that of the current POWER3 SMP Thin and Wide Nodes), a range of PCI adapters, and include integrated 10/100 Mbps Ethernet and Ultra2 SCSI interfaces.

Table 91 outlines the potential memory DIMMS available.

Table 91. Minimum and Maximum Memory

Feature Code	Description	Minimum	Maximum
4098 - default	MemoryExpansion Feature	2	2
4110 - default (1)	256 MB (2 x 128 MB) SDRAM DIMM	1	16
4119	512 MB (2 X 256 MB DIMM)	0	16
4133	1 GB Memory (8 X 128 MB DIMM)	0	4
4402	2 GB Memory (16 X 128MB DIMM)	0	2
4403	3 GB Memory (24 X 128 MB DIMM)	0	1
4404	4 GB Memory (32 X 128 MB DIMM)	0	1
4412	4 GB Memory (16 X 256 MB DIMM)	0	2
4413	6 GB Memory (24 X 256 MB DIMM)	0	1
4414	8 GB Memory (32 X 256 MB DIMM)	0	1

Disk Requirements and Options

The 375 POWER3-II SMP Thin Node can support up to 72.8GB of mirrored internal disk. The option of ordering no internal disks is available for customers who choose to use external disk storage. This is done by not selecting any disk pairs offered and selecting (# 9121) for external SSA boot or (# 9124) for external SCSI Boot. Configuring these nodes for mirroring is strongly recommended.

Table 92 outlines the potential disks, which only come in pairs, that can be installed on the 375 MHz POWER3-II Thin Node.

Table 92. Potential Disk Pairs in 375 MHz Thin Node

Feature Code	Descriptions
2909	9.1 GB Ultra SCSI DISK Drive Pair

Feature Code	Descriptions
2918	18.2 GB Ultra SCSI Disk Drive Pair
3804	9.1 GB 10K RPM Ultra SCSI Disk Pair
3810	18.2 GB 0K RPM Ultra SCSI Disk Pair

Switch Adapter Requirements

POWER3-II 375 MHz SMP Nodes are connected to the SP Switch using the SP Switch MX2 Adapter (#4023), which attaches directly to the processor mezzanine (MX) bus, providing increased bandwidth for TCP/IP, UCP/IP, and user-space communications between nodes. It does not require a PCI slot, which leaves one more adapter slot available for application use.

In a system configured with an SP-8 switch, any MES that increases the form factor of a node, such as thin to wide node or wide to high node, will require *all* nodes that are above the node being upgraded to be deleted via the `spdelnode` command and reinstalled. Additionally, the SPS switch cables on *all* the nodes above the one being upgraded need to be recabled to represent the new switch node positions.

Dynamic Processor Deallocation

Another high availability feature is dynamic processor deallocation. This feature allows the system to completely isolate a processor that has been determined to be operating below an acceptable level of reliability. To accomplish this, the system monitors processor status in real-time to determine when the number of recoverable errors has exceeded a specified threshold. When the threshold is exceeded, the system sends a report to the operating environment that a processor is unreliable. The AIX operating environment then re-routes all interrupts and processes to other processors and disables the unreliable processor complex without interrupting end user services. Firmware level 991216 must be installed on the 375 MHz Thin Node. Since it is a feature that was introduced to AIX in February 7, 2000, the AIX section of this book discusses this function further.

The location of the URL that microcode or firmware can be downloaded from is:

<http://www.austin.ibm.com/support/micro/download.html>

7.6.6 332 MHz SMP Thin Nodes

332 MHz SMP thin nodes (# 2050) have PCI bus architecture and use either two or four 332 MHz PowerPC processors per node. These nodes are functionally equivalent to an IBM RS/6000 7025-F50 workstation. The SP system must be operating at PSSP 2.4 (or later) to use these nodes.

The 332 MHz SMP thin node occupies one half of a drawer and may be installed singly with systems operating at PSSP 3.1 or later. Therefore, up to 16 SMP thin nodes can be housed in a tall frame. When installed singly, 332 MHz SMP thin nodes must be placed in the odd numbered node slot. Uniprocessor thin nodes are still required to be installed in matched pairs that fill the node drawer.

For electromagnetic compliance, these nodes are housed in an SMP Enclosure. This enclosure (# 9930) is automatically included when you order a 332 MHz SMP thin node. For installations using single SMP thin nodes, a cover plate (# 9931) is also included to cover the unused enclosure slot.

If you are going to mount a 332 MHz SMP thin node into an older 2.01 m or 1.25 m frame, a power system upgrade is necessary. However, once you have done the power system upgrade, these nodes are fully compatible with all existing SP hardware except for High Performance Switches.

Note

332 MHz SMP thin nodes are not compatible with High Performance Switches (# 4010 and # 4007).

Bus Description

The 332 MHz SMP thin node PCI bus contains two 32-bit slots PCI slots (slots I2 and I3). The I1 slot is reserved for the optional SP Switch MX2 Adapter.

Previously installed 332 MHz SMP thin nodes may have a withdrawn SP Switch MX Adapter in the I1 slot.

Single SMP Thin Node Configuration Rules

With PSSP 3.1, single POWER3 SMP thin nodes and single 332 MHz SMP thin nodes are allowed in both tall and short frame configurations provided the following rules are observed:

- Single SMP thin nodes must be installed in the odd numbered node position.

- Single SMP thin nodes are not supported in the even numbered node position.
- Empty node drawers are allowed on tall frames if the frame is either a non-switched frame or configured with an SP Switch (16-port switch).
- Tall frames configured with the SP Switch-8 (8-port switch) must have all nodes placed in sequential order; no empty drawers are allowed. Therefore, the single SMP thin node in these frames is the last node in the configuration.
- Short frame configurations must have all nodes placed in sequential order; no empty drawers are allowed. Therefore, the single SMP thin node in these frames is the last node in the configuration.
- A single POWER3 SMP thin node and a single 332 MHz SMP thin node each occupy one half of a node drawer.
- Single POWER3 SMP thin nodes and single 332 MHz SMP thin nodes may be mixed in a thin node drawer.
- If a frame has more than six single SMP thin nodes installed, that frame will have an uneven weight distribution. You must be careful when moving these frames.
- Uniprocessor thin nodes must still be installed in matched pairs and occupy a full node drawer.

Mandatory Prerequisites

332 MHz SMP thin nodes occupy one half of a node drawer. When two SMP thin nodes are placed in one drawer, the nodes may be asymmetrically configured for memory, DASD, processor speed, and adapters. Up to 16 332 MHz SMP thin nodes may be installed in one tall frame and up to eight in a short frame.

The mandatory prerequisites are:

- PSSP 2.4 (or later) on the Control Workstation, backup nodes, and processor node
- Two processors (mounted in one slot)
- 256 MB of memory
- 9.1 GB of DASD (with internal booting)
- An upgraded power system on older frames

Standard Features

Each 332 MHz SMP thin node is functionally equivalent to an RS/6000 7025-F50 and has:

- Two processor slots allowing a maximum of four processors per node
- Two memory slots supporting up to 3 GB of memory
- Two DASD bays supporting up to 36.4 GB of storage
- A dedicated mezzanine (MX) bus slot for an optional switch adapter
- Two 32-bit PCI slots for communication adapters
- Integrated 10BaseT/10Base2 Ethernet (only one port may be used at one time)
- Integrated SCSI-2 Fast/Wide
- Standard Service Processor
- External nine-pin RS-232 on the planar S2 port
 - This connection has active heartbeat and is available for customer applications.

Processor Requirements and Options

SMP thin nodes require a minimum of two 332 MHz PowerPC processors mounted on one card. However, you can order an additional processor card (# 4320) to configure the node with a total of four CPUs. Table 93 provides the processor options for the SMP nodes.

Table 93. Processor Options for 332 MHz SMP Thin Nodes

Feature Code	Multiplier	Description	Comments
4320	X 1	One processor card with two CPUs	Minimum required
4320	X 2	Two processor cards with two CPUs each (four CPUs total)	Maximum allowed

Memory Requirements and Options

332 MHz SMP thin nodes have two memory cards and require a minimum of 256 MB of memory. These nodes support a maximum of 3 GB of memory. Memory is supplied by 128 MB DIMMs that must be mounted in pairs (256 MB increments).

The memory cards are not required to be configured symmetrically. Each card has the capacity to mount 2 GB of DIMMs; however, only 3 GB are addressable per node. Memory cards and DIMMs are not interchangeable

between SMP and non-SMP thin nodes. Table 94 provides the available memory features for the SMP thin nodes.

Table 94. Memory Features for 332 MHz SMP Thin Nodes

Feature Code	Description	Minimum Node Requirement	Maximum Allowed Per Node
4093	Base Memory Card	2	2
4110	One Pair of 128 MB DIMMs (256 MB total)	1 pair	6 pairs
4133	Eight 128 MB DIMMs (1 GB total)	0	2
4402	16 128 MB DIMMs (2 GB total)	0	1
4403	24 128 MB DIMMS (3 GB total)	0	1

Disk Requirements and Options

332 MHz SMP thin nodes can have up to two internal DASD attached through an integrated SCSI-2 network. The 332 MHz SMP thin node can have either no internal DASD (with external booting) or from 9.1 GB to a maximum of 36.4 GB of internal disk storage. External storage devices can be accessed through an optional Ultra SCSI Adapter (# 6207) or SCSI-2 Adapter (# 6209).

Optional direct access storage devices are available as follows:

- 9.1 GB Ultra SCSI disk drive pair (# 2909)
- 18.2 GB Ultra SCSI disk drive pair (# 2918)

This node does not require special cables or adapters to mount an internal disk.

Switch Adapter Requirements

The switch adapter for SMP thin nodes does not occupy a PCI slot. Instead, the switch adapter for these nodes is installed into the Mezzanine (MX) bus. The MX bus connects the I/O planar with the system planar. Placing the switch adapter in the MX bus enables switch traffic to proceed at higher bandwidths and lower latencies.

In switch configured systems, 332 MHz SMP thin nodes require the following switch adapter:

SP Switch MX2 Adapter (# 4023)

332 MHz SMP Thin Node Switch Restrictions

The 332 MHz SMP thin node is not compatible with the older High Performance series of switches.

If an SMP thin node is going to be placed into an SP system configured with a switch, that switch must be either an SP Switch or an SPS-8 switch.

Switch adapters for SMP thin nodes are not interchangeable with switch adapters used on uniprocessor thin nodes.

SMP Thin Nodes in Expansion Frames

With PSSP 3.1 and single SMP thin nodes, the restriction on using SMP thin nodes in expansion frames has been partially removed. The following applies to single SMP thin nodes only:

- Frames with single SMP thin nodes installed can be used as nonswitched expansion frames. Similarly, if a frame has single SMP thin nodes installed and a switch with unused switch ports, it can have a nonswitched expansion frame attached to the unused switch ports.
- However, if a frame has a node drawer containing a pair of SMP thin nodes, the frame cannot be used as an expansion frame and cannot have expansion frames attached.
- Single SMP thin nodes must be placed in the odd numbered node slot.
- One or more single SMP thin nodes can be installed in frames with the 16-port SP Switch.
- Frames with uniprocessor thin nodes or SMP thin node pairs installed require an SP Switch for expansion.

7.6.6.1 Single SMP Thin Node Configuration Rules

Beginning with PSSP 3.1, single 200 MHz POWER3 SMP thin nodes and single 332 MHz SMP thin nodes are allowed in both tall and short frame configurations provided the following rules are observed:

- Single SMP thin nodes must be installed in the odd numbered node position. They are not supported in the even numbered node position.
- Empty node drawers are allowed on tall frames if the frame is either a non-switched frame or configured with an SP Switch (16-port switch).
- Tall frames configured with the SP Switch-8 (8-port switch) must have all nodes placed in sequential order; no empty drawers are allowed. Therefore, the single SMP thin node in these frames is the last node in the configuration.

- Short frame configurations must have all nodes placed in sequential order; no empty drawers are allowed. Therefore, the single SMP thin node in these frames is the last node in the configuration.
- A single POWER3 SMP thin node and a single 332 MHz SMP thin node each occupy one half of a node drawer.
- Single POWER3 SMP thin nodes and single 332 MHz SMP thin nodes may be mixed in a thin node drawer.
- If a frame has more than six single SMP thin nodes installed, that frame will have an uneven weight distribution. You must be careful when moving these frames.
- Uniprocessor thin nodes must be installed in matched pairs and occupy a full node drawer.

7.7 The SP Switch

The key element of the SP communication subsystem is the SP Switch.

The design objectives for the SP Switch network is low-latency, scalable high-bandwidth, and fault-tolerant communication to support the demands of parallel and client-server applications.

The SP Switch provides the message passing network that connects all nodes with a minimum of four paths between any pair of nodes.

7.7.1 SP Switch Components

This section discusses hardware components of the switch used in the SP system.

7.7.1.1 SP Switch Link

An SP Switch link connects two network devices. It contains two channels carrying packets in opposite directions. Each channel includes 8 bits Data, 1 bit Data Valid, and 1 bit Token signal.

7.7.1.2 SP Switch Ports

An SP Switch port is part of a network device and is connected to other SP Switch ports through the SP Switch link. The SP Switch port includes two ports (input and output ports) for full duplex communication. The following diagram shows the relationship between the SP Switch link and the SP Switch port. Figure 60 on page 287 shows the data flow between switch ports.

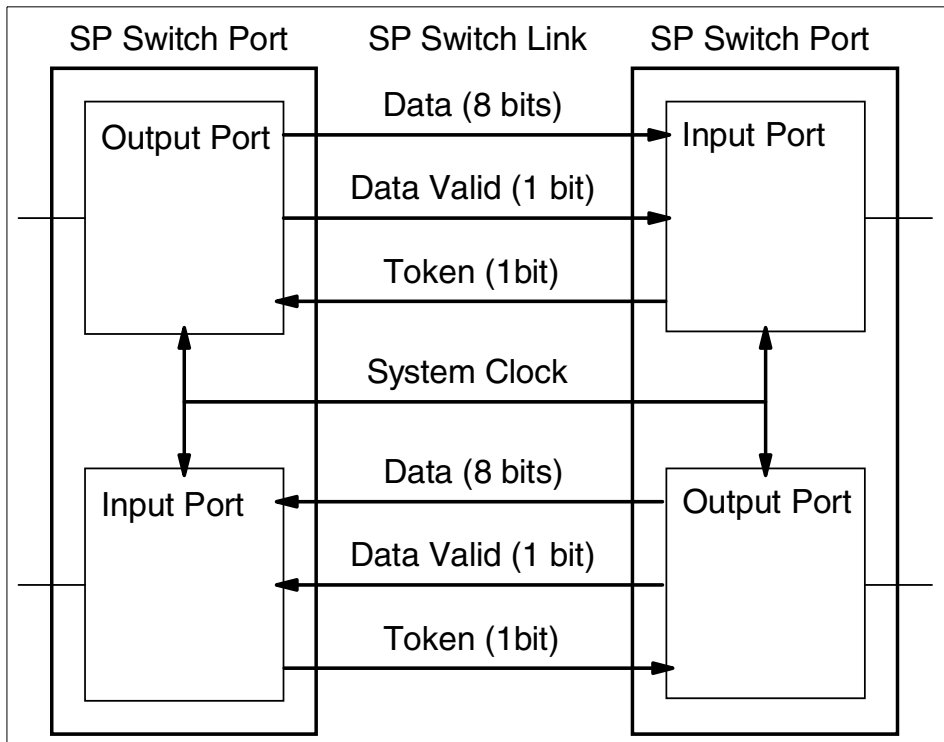


Figure 60. SP Switch Link and SP Switch Port

7.7.1.3 SP Switch Chip

An SP Switch chip contains eight SP Switch ports and a crossbar that allows packets to pass directly from port to port. These crossbar paths allow packets to pass through the SP Switch with low latency. As soon as an input port decodes the routing information carried by an incoming packet, it sends a crossbar request to the appropriate output port.

Figure 61 shows the SP Switch ports, the crossbar, and the SP Switch chip.

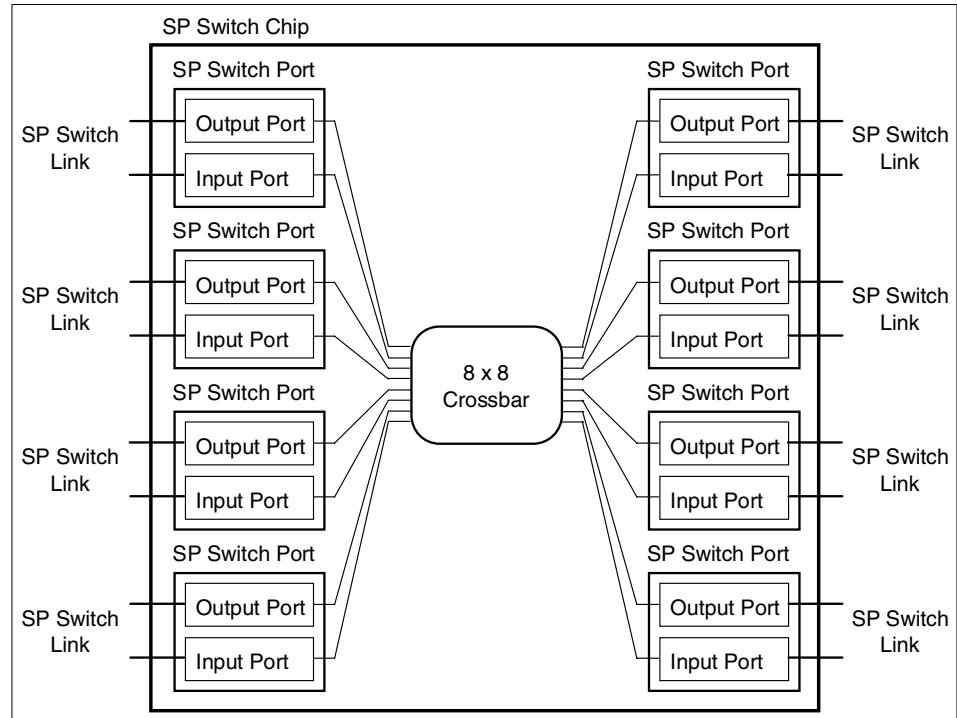


Figure 61. SP Switch Chip

7.7.1.4 SP Switch Board

An SP Switch board contains eight SP Switch chips that provide connection points for each of the nodes to the SP Switch network as well as for each of the SP Switch boards to the other SP Switch boards. There are 32 SP Switch ports in total. Of these, 16 could be connected to nodes and the other 16 to other SP Switch boards. The SP Switch board is mounted in the SP Frame.

Figure 62 shows the interconnections used on the switch board.

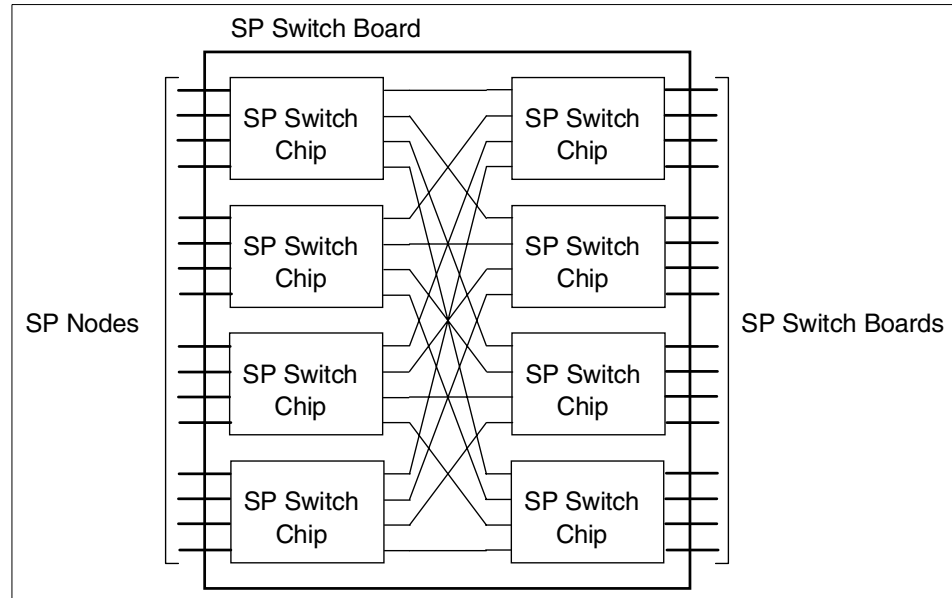


Figure 62. SP Switch Board

7.7.1.5 SP Switch Adapter

An SP Switch Adapter includes one SP Switch port that is connected to an SP Switch board. The SP Switch Adapter is installed on an SP node.

An SP Switch Adapter is designed for nodes that have both a PCI bus and an MX or MX2 bus.

Figure 63 provides a logical diagram of an SP Switch Adapter.

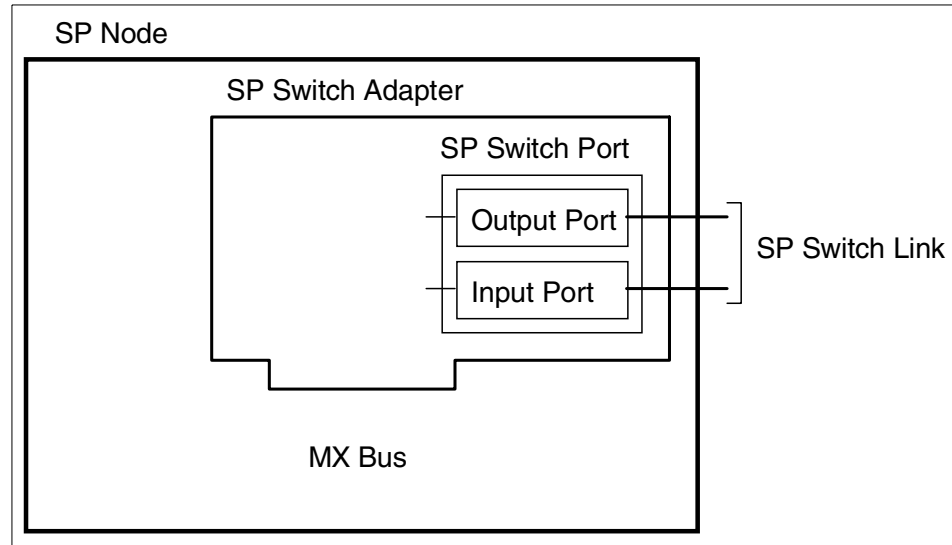


Figure 63. SP Switch Adapter

7.7.1.6 SP Switch System

Figure 64 shows the SP Switch system. In one SP frame, there are 16 nodes equipped with SP Switch Adapters and one SP Switch board. Sixteen SP Switch Adapters are connected to 16 of 32 SP Switch ports in the SP Switch board. Another 16 SP Switch ports are available for other SP Switch boards.

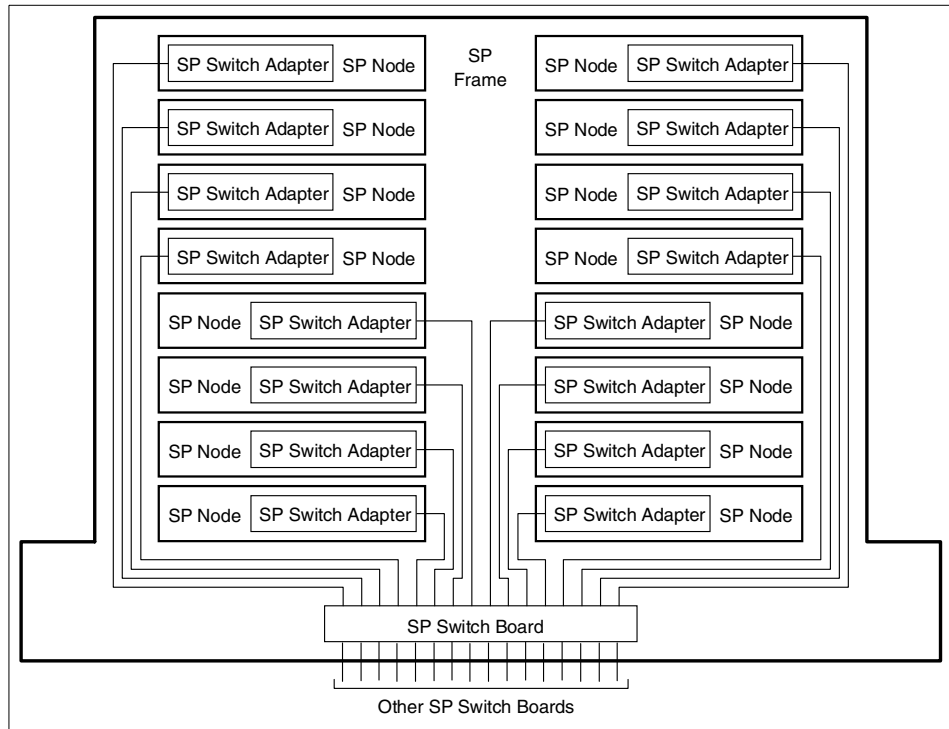


Figure 64. SP Switch System

7.7.2 SP Switch Types

This section describes the switch types installed on SP systems

7.7.2.1 SP Switch (# 4011)

The most current switch type is simply called the SP Switch (SPS). It offers the following improvements over the older High Performance Switch (HPS):

- Higher availability
- Fault isolation
- Concurrent maintenance of nodes
- Improved switch chip bandwidth

SP Switches use the SP Switch MX2 Adapter (# 4023) installed in nodes with a PCI I/O bus. The SP Switch MX2 Adapter now replaces the older SP Switch MX Adapter (# 4022). One switch adapter is needed for each node in the SP system. The required SP Switch Adapter or SP Switch MX2 Adapter connects each SP node to the SP Switch board.

7.7.2.2 High Performance Switch (# 4010)

The High Performance Switch (HPS) is not compatible with the SP Switch. High Performance Switches cannot be mixed with SP Switches in a system or in separate system partitions. To take advantage of the SP Switch's performance, all High Performance Switches must be upgraded to SP Switches.

The High Performance Switch is being phased out and is not available for new systems; however, they are still available for existing systems that are already equipped with High Performance Switches.

SP Switch-8 (# 4008) and High Performance LC-8 Switch (# 4007)

Eight-port switches are a low cost alternative to the full size 16-port switches. Both the 8-port SP Switch-8 (SPS-8) and the 8-port High Performance Switch (HiPS LC-8) provide switch functions for an 8-node SP system. The SP Switch-8 is compatible with high nodes, but the HiPS-LC8 is not.

The SP Switch-8 is the only low cost switch available for new systems; however, the High Performance LC-8 is still available for existing systems.

The SP Switch-8 has two active switch chip entry points. Therefore, the ability to create system partitions is restricted with this switch. With the maximum eight nodes attached to the switch, there are two possible system configurations:

- A single partition containing all eight nodes

- Two system partitions containing four nodes each

The High Performance LC-8 only has one switch chip. Therefore, only a single partition is possible.

7.8 Control Workstation

An advantage of an SP is that it has a single point of control. This is the Control Workstation. The Control Workstation is a separate RS/6000 with a color graphics monitor that serves as point-of-control for managing and maintaining the SP nodes.

The workstation connects to the frame by using a 15 m IBM RS232 line, which is supplied with each frame, and by using another form of network: Ethernet. A system administrator can log into the Control Workstation from any workstation on the network to perform system management, monitoring, and control tasks. These two different types of connections are necessary because the serial link is used for hardware-related control and for monitoring communications, and the Ethernet is used for system management commands, which are AIX based. The serial connection is also used for hardware-related communications even before AIX has started up on the node.

The Control Workstation also acts as a boot/install server for other servers in the SP system using an AIX utility called Network Installation Manager (NIM). In addition, the Control Workstation can be set up as an authentication server using Kerberos. It can be the Kerberos primary server, with the master database and the administration service as well as the ticket-granting service. Or it can be set up as a Kerberos secondary server, with a backup database, or just the ticket-granting service. The function and connectivity of a Control Workstation is shown in Figure 65 on page 294.

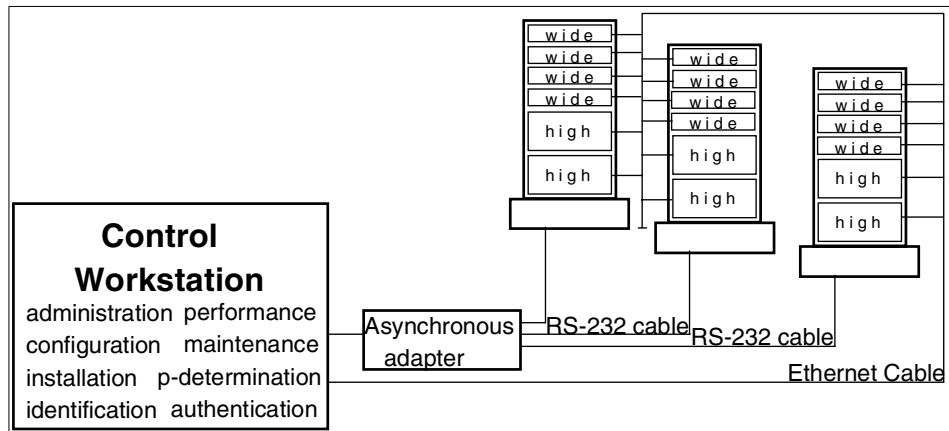


Figure 65. Control Workstation Connectivity and Function

Some of the newer Control Workstations, such as the 7025-F50, offer a support processor as either a standard or an optional feature. The support processor is a standard processor that handles system startup with some system monitoring functions. When this option is installed with a modem on the first serial port, remote system restarts can be performed on SPs located in unmanned locations.

7.8.1 High Availability Control Workstation

A high availability option is available for Control Workstations that will enable a second Control Workstation to be connected. This option eliminates the Control Workstation as the single point of failure. When the primary Control Workstation becomes unavailable, either through a planned event or a hardware or software failure, the SP high availability component detects the loss and shifts that component's workload to the backup system.

To provide this extra availability and to eliminate the Control Workstation as a single point of failure, additional software and hardware are needed, such as HACWS and its necessary pre-requisites.

7.8.2 Supported and Unsupported SP Control Workstations

The Control Workstation connects to the SP system both through the Ethernet and through the 15 m RS-232 cable. System management commands and monitoring go through the Ethernet, and low-level hardware management information goes through the RS-232 cable. Table 95 on page

295 lists the various models of RS/6000 that are supported and unsupported for use as an SP Control Workstation.

Table 95. Supported and Unsupported SP Control Workstations

Machine	Model	Notes	PSSP 3.1 Support
7006	41X	Only if HACWS is never going to be configured.	N
	42X	Only if HACWS is never going to be configured (or a new 3.2.5 partition is not required).	N
7007	N40		N
7008	XXX		N
7009	CXX	Not recommended beyond one frame due to slow tty response.	N
7011	22X, M2X, 23X		N
	25X	Not recommended, total system performance is slow.	N
7012	3XX (< 370)		N
	3XX (≥ 370)		Y
	GXX		Y
7013	5XX (< 570)		N
	5XX (≥ 570)		Y
	JXX	Requires a 7010 Model 150 X-Station and display. Other models and manufacturers that meet or exceed this model can be used. An ASCII terminal is required as the console.	Y
7015	9XX (< 970)		N

Note: Shaded areas represent unsupported options.

Machine	Model	Notes	PSSP 3.1 Support
	9XX (>= 970) RXX	Requires a 7010 Model 150 X-Station and display. Other models and manufacturers that meet or exceed this model can be used. An ASCII terminal is required as the console.	Y
7020	40P		N
7024	E20, E30	PSSP 2.2 and above.	Y
7025	F30	PSSP 2.2 and above. New systems require either 8-port (#2493) or 128-port (#2944) PCI bus asynchronous adapters for frame controller connections.	Y
	F40	PSSP 2.2 and above. Cannot use native serial ports. Requires 8-port (#2943) or 128-port (#2944) PCI bus asynchronous adapter for frame controller connections. These adapters require AIX 4.2.1 or 4.3 on the CWS.	Y
	F50	PSSP 2.2 and above. Cannot use native serial ports. Requires 8-port (#2943) or 128-port (#2944) PCI bus asynchronous adapter for frame controller connections. These adapters require AIX 4.2.1 or 4.3 on the CWS.	Y
7026	H10, H50	PSSP 2.2 and above. Cannot use native serial ports. Requires 8-port (#2943) or 128-port (#2944) PCI bus asynchronous adapter for frame controller connections. These adapters require AIX 4.2.1 or 4.3 on the CWS.	Y
7030	3XX		Y
7043	140, 240	PSSP 2.2 and above. Cannot use native serial ports. Requires 8-port (#2943) or 128-port (#2944) PCI bus asynchronous adapter for frame controller connections. These adapters require AIX 4.2.1 or 4.3 on the CWS. Only supported with up to four frames (no matter the number of nodes in these four frames).	Y
7247	82X		N
Note: Shaded areas represent unsupported options.			

Machine	Model	Notes	PSSP 3.1 Support
7248	1XX Model 43P		N
7249	851		N
7586	43P		N
Note: Shaded areas represent unsupported options.			

At the time of publication, the Model 7043-140 and the 7025-F50 are the only marketed Control Workstations.

Note

The 7043 cannot be used for SP systems with SP-attached servers such as Models S70 Advanced or the Model S80. Read the following section for further details.

The 7043 can only be used on SP systems with up to four frames. This limitation applies to the number of frames and not the number of nodes. This number includes expansion frames. The 7043 cannot be used for SP systems with SP-attached servers, such as RS/6000 Enterprise Servers Models S80 or S70 Advanced. Each S80 is considered a frame as well as a node. That would mean that two Model S80s and two regular node frames would be the maximum that the 43P could support.

If the Model S80 is a *node* attached to an SP, the *Planning Guide* states that a Model 140 cannot be the CWS - the Model F50 must be the CWS. This is because the Model S80 uses more CPU resources than other nodes. Typically, the Model F50 is configured with 256 MB minimum and four 9.1 GB disks as a minimum for two frame systems. It is recommended that an additional 256 MB is added for the next two SP frames, and then up to 1 GB for larger systems.

7.8.3 CWS Disk and Memory Requirements

The minimum requirements for the Control Workstation are:

- At least 128 MB of main memory. An extra 64 MB of memory should be added for each additional system partition. For SP systems with more than 80 nodes, 256 MB is required, and 512 MB of memory is recommended.
- Four GB of disk storage. If the SP is going to use an HACWS configuration, you can configure 2 GB of disk storage in the rootvg volume group and 2 GB in an external volume group.

Because the control workstation is used as a Network Installation Manager (NIM) server, the number of unique file sets required for all the nodes in the SP system might be larger than a normal single system. You should plan to reserve 2 GB of disk storage for the file sets, and 2 GB for the operating system. This will allow adequate space for future maintenance, system mksysb images, and LPP growth. Keep in mind that if you have nodes at different levels of PSSP or AIX, each node requires its own LPP source which will take up extra space.

A good rule of thumb to use for disk planning for a production system is 4 GB for the rootvg to accommodate additional logging and /tmp space, plus 1 GB for each AIX release and modification level for lppsource files. Additional disk space should be added for mksysb images for the nodes.

If you plan on using rootvg mirroring, then for one mirror, double the number of physical disks you estimated so far. For two mirrors, triple the estimate.

The most current information about the CWS is available using the following resources:

- *Volume 2, Control Workstation and Software Environment, GA22-7281*
- http://www.rs6000.ibm.com/resource/aix_resource/sp_books/hardware/index.html

Information about the supported Control Workstations can be found at:

- <http://www.s390.ibm.com:80/bookmgr-cgi/bookmgr.cmd/books/sspsp311/1%2e2%2e10%2e1%2e3?SHELF=>

7.9 The SP Switch Routers

The RS/6000 SP Switch Routers (9077-04S and 9077-16S) are high-performance I/O gateways for the SP system and provide the fastest available means of communication between the SP system and the outside world, or among multiple SP systems. Figure 66 on page 299 shows the switch router's role as a gateway.

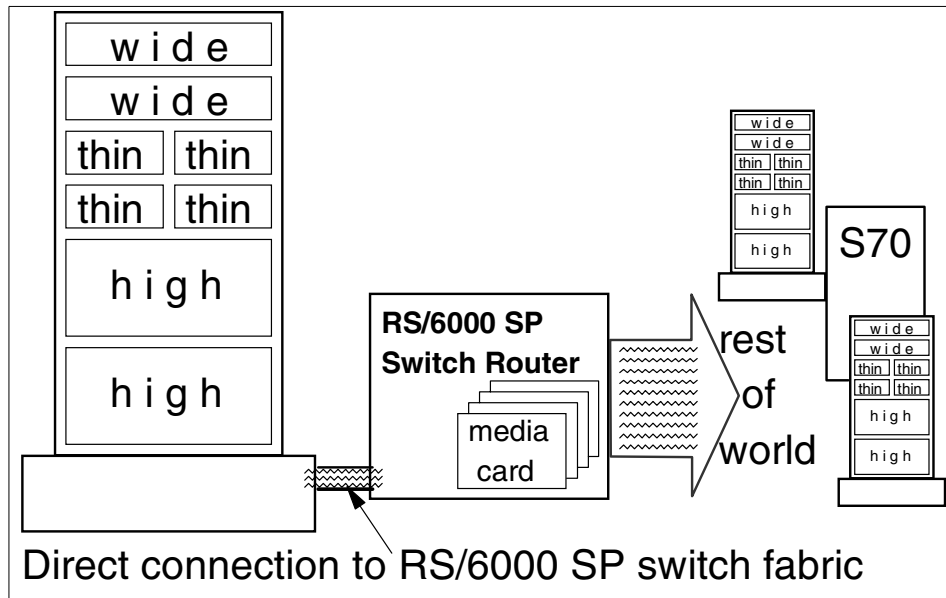


Figure 66. RS/6000 SP Switch Router Offers Many Possibilities

The RS/6000 SP Switch Router is based on the Ascend GRF switched IP router product from Ascend Communications, Inc. You can order the SP Switch Router as the IBM 9077. This enables direct network attachment to the SP Switch. The 9077-04S has four slots for media cards, and the 9077-16S has 16. These provide network attachment. Both models come with an SP Switch Router Adapter in one slot to provide attachment to an SP Switch port. The remaining slots can be used by other media cards to provide attachment to external networks. These slots can also be used by additional SP Switch Router Adapters to provide scalable, incremental steps of SP I/O performance for a single SP system, alternate paths to an SP system, and interconnection for multiple SP partitions and systems. Each media card has its own IP engine, enabling SP I/O to scale nearly linearly with the number of cards.

7.9.1 Gigabit Ethernet Media Cards for the SP Switch Router

The SP Switch Router is a high-performance I/O gateway for the RS/6000 SP system, which provides the fastest available means of communication between the SP system and the outside world, or among multiple SP systems. It connects the SP Switch to various popular IP networks, including 10/100 Ethernet, ATM, HIPPI, and FDDI. Now, this list of networks is being

expanded to allow high-speed SP interconnection directly with Gigabit Ethernet networks.

Two new Gigabit Ethernet media cards are available for the SP Switch Router:

- 1000BaseSX media card (for short distances of up to 550 meters)
- 1000BaseLX media card (for long distances of up to 5 kilometers.)

Both media cards support the standard 1518 byte Ethernet packets and the 9000 byte Jumbo Frames.

7.9.2 The Advantages of a Switch Router

Although an SP node can be equipped with a variety of network adapters and used to make network connections for the whole system, the RS/6000 SP Switch Router with the Switch Router Adapter and optional network media cards offers many advantages when connecting to the external networks, as follows:

- It offers performance to system users who need faster I/O than that provided by node I/O adapters. This includes research organizations, universities, and others relying on leading-edge technical computing applications that need to quickly move large amounts of data between the SP system and external networks and commercial computing installations that need high-speed access to data servers, such as the RS/6000 S80.
- It offers availability. A range of features designed for use in high-availability environments can improve the system up-time, Dynamic network configuration bypasses, failed paths using standard IP protocols; media cards are hot-swappable and can be removed or added while the system is running; Power supplies are redundant, load-balancing, and hot-swappable; the SP Switch Router can be accessed and rebooted remotely; and multiple SP Switch Routers can be connected to the SP system to provide alternate paths in case of failure of one. Support for industry-standard features allows administrators to monitor attached networks with existing management tools. For example, the SP Switch Router works with standard Simple Network Management Protocol (SNMP) management packages.
- It offers price performance. The SP Switch Router can cut communication costs of the SP system and offer better price/performance than dedicated SP I/O nodes. Since the I/O gateway role can now be performed by the SP Switch Router, the cost of additional processor nodes and I/O adapters can be avoided. In many cases, this lowers the overall cost of an I/O gateway. Also, since SP I/O nodes typically connect to an external hub or

router, this cost can be eliminated. And, in installations where there are multiple RS/6000 SP systems or partitions, a single SP Switch Router can act as a common gateway, reducing the gateway cost for each additional SP system or partition to just the cost of an SP Switch Router Adapter. This is especially useful where multiple SP systems or partitions need to interconnect.

7.10 SP Attached Servers - The S-Series

Since October 5, 1998, for the Model S70 and S70 Advanced, and November 26, 1999, for the Model S80, the S-Series RS/6000 is able to function as an attached SMP server to a switched or non-switched tall frame configuration, enabling the SP control workstation to be a single point of control within the IBM RS/6000 SP environment and operating under the control of the IBM Parallel Systems Support Programs for AIX. This interconnection can be accomplished utilizing the IBM RS/6000 SP System Attachment Adapter (# 8396) or via an Ethernet connection. Some I/O adapters available on the Model S80 systems are not supported in the SP environment and must be removed. Refer to Appendix D, "Adapter Placement Guide" on page 661.

When attached to an SP system, the S80 or S7A, although external to the SP frames, is managed by the system as an SP node and can function from both an application and management standpoint. IBM Parallel System Support Programs for AIX, Version 3.1 (5765-D51), or later, is required for a Model S80 system to function as an attached SMP server within the IBM RS/6000 environment.

A minimum of one Ethernet adapter is required for a Model S80 or S7A system to function as an attached SMP server within the IBM RS/6000 SP environment. This adapter must be recognized by the Model S80 or S7A as en0 and must reside in slot #5 of the primary I/O drawer.

7.10.1 Floor Placement Considerations

Floor location of the elements of SP-attachment are critical elements of the pre-attachment planning process.

The Control Workstation to S7A/S80 connection cables are 15 meters in length. The SP Switch to S7A/S80 connection cable is 10 meters in length. The effective lengths of these cables are somewhat less than this after subtracting for vertical distances to and from the under-floor area. That means you must do some floor placement planning to account for usable switch cable length (10 meters, 6.5 effective meters distance between the S7A/S80 and the SP Switch) and RS-232 cable lengths (15 meters, 11.5

meters effective cable length distance between the S7A/S80 and the Control Workstation).

It is possible to attach up to 16 S7A/S80 servers around a single SP frame, but planning is necessary. Planning must also include work with the *IBM General Information Installation Manual - Physical Planning, GC22-7072*. This document will help you plan for space requirements, site environment considerations, air conditioning, electrical power and grounding, cabling, and communications (since your Control Workstation will be phoning home using a modem if trouble develops).

7.10.2 HACWS Considerations with SP-Attached Servers

The High Availability Control Workstation (HACWS) requires special consideration with SP-attached servers since hardware control failover will not succeed without two manual intervention steps. The customer must determine if HACWS will meet their needs under these restrictions. Note that despite the limitations of HACWS support with the SP-attached server, customer jobs executing on the SP-attached server are unaffected during the failover and hardmon daemon control recovery procedure. Only hardware control of the SP-attached server is affected.

The SP-attached server has two serial connections between the Control Workstation and the S7A/S80. These SP-attached connections provide S1 Term capability and hardmon controls, respectively, to the SP-attached server. At this time, the ability to split or y-connect these connections to provide attachment to two Control Workstations and automated failover does not exist.

Customers may wish to work around this product characteristic in one of two ways:

- Alternate 1: Use a 2-in/4-out A-B switch to provide manual switching of Control Workstation attachment once failover has occurred. Recycle (down and up) the hardmon daemon to restore hardware control.
- Alternate 2: Physically unplug cables between one Control Workstation/SP-attached server and plug in cables between the second Control Workstation and the SP-attached server. Recycle (down then up) the hardmon daemon to restore full hardware control.

Note

- The preceding process is not documented in any official publications for this release.
- The SP administrative LAN Ethernet is available from the backup CWS after the failover.
- Switch communications to the SP-attached server are not affected during CWS failover.
- Running jobs on the SP-attached server are not impacted by the CWS failover.

7.10.2.1 Maximum Number of SP-Attached Servers per SP System

The maximum number of SP-attached servers allowed to be attached to a single SP system is 16. That means you must have up to 16 available switch connections on your SP system (one per SP-attached server). If your frame does not have enough available switch connections to support the number of switch-attached servers you desire, you must purchase an additional frame and switch to support the SP attachment.

7.10.2.2 Switchless Attachment

If you have chosen to have a switchless attachment to your S70 Advanced/S80, your SP system must also be switchless. You must have available (unused) switch node numbers in your System Data Repository even though you are not using a switch. That means you cannot attach an S7A/S80 to a single frame system where all 16 nodes reside.

7.10.2.3 One Attachment Adapter per S7A/S80

Only one RS/6000 SP System Attachment Adapter is permitted in each SP-attached S-Series server. This is consistent with the rules for SP nodes. More SP System Attachment Adapters in a single SP-attached server are not supported by PSSP 3.1 code.

7.10.2.4 PCI-Attachment Exclusive to the S7A/S80

The RS/6000 SP System Attachment Adapter is not supported (and will not work) in any other PCI-bus equipped device.

7.10.2.5 Separate PSSP 3.1 License Required

You must order a separately chargeable PSSP 3.1 license against each SP-attached server's serial number. This is true for each Licensed Program Product your customer plans to install on the SP-attached server. You cannot use the licenses your customer has secured for their SP system. Even though

this S80 is attached to an SP system, runs PSSP software, and is transparently managed by the Control Workstation as any SP node would be, it is still a separate entity and so requires separate licensing.

7.10.3 Pre-Install Required Meeting - System Assurance Reviews

Before the order can be placed and the hardware installed, there will be a required Solutions Assurance Review (SAR) between Marketing, CEs, and the customer to ensure all participants understand prerequisite planning and activities for the SP-attach activities.

7.10.3.1 Required Pre-SP-Attachment Activities

There are a number of potential system outage activities that must take place before the actual SP-attachment is made. These include both limited outage activities and system-wide outage activities. Limited system outage activities might include those activities that impact part of a system, such as possible movement or removal of adapters within the S7A/S80 or a planned failover between S7A/S80 stand-alone systems to permit the installation of code upgrades on those stand-alones. System-wide outage activities might include the installation of new code levels on the Control Workstation and the subsequent required reboot. Planning for all of these activities will come out of the SAR meeting, and all should be conveyed to the customer as a necessary prerequisite to the merging of the S7A/S80 into the SPs environment.

7.10.4 64-bit Application Tolerance

While the S7A and S80 are 64-bit machines and fully support 64-bit programs, the SP-attached server's PSSP 3.1 code will provide 64-bit tolerance. That is to say, 64-bit applications and programs will be supported as long as they do not call upon any PSSP program services (APIs). Customers may assume SP support for 64-bit programs unless this point is explicitly covered in the SAR or in direct conversation. For a detailed discussion of this topic, refer to a white paper available at:

www.rs6000.ibm.com/resource/technology/#sp

7.10.5 Service

To ensure that a good service response is provided, the following elements need to be considered.

7.10.5.1 When a Service Call Is Required

Service is handled by either RS/6000 service or SP service organizations based on the first suspected faulty component (SP or S7A/S80). If at any

time, the handling service organization perceives that this should be in the other service flow, there is a soft hand-off to avoid any delay in servicing the customer's problem. With this system of parallel call flows with soft hand-off, a customer or Customer Engineer should never feel that they have ended up in the wrong queue.

7.10.5.2 Service Director Phones Home for the Customer

The IBM Service Director should be activated on the customer's machines so that when there is a S7A/S80 or SP hardware problem, service will be called and the problem will be logged automatically. If the Service Director is properly configured, calls home will send 7017-generated calls to S80 support/CEs, and 9076SP-generated calls will be sent to SP trained CE/support. This package includes a modem.

7.10.6 Ordering Through Joint Configurator Sessions

Ordering the SP-attached server requires the ordering of both SP features and RS/6000 features. Care needs to be taken when ordering any adapters for the S80 or S70 Advanced to make sure they are supported in the SP operating environment. Ordering the SP-attach features to join the S7A/S80 with the SP environment means that careful preplanning through the SAR meeting has already been completed. Adapter moves and S7A/S80 boot device redefinitions have been scheduled or executed during scheduled system outage time, and software upgrades on the Control Workstation have been ordered and scheduled. It is now time to order the SP-attach software and hardware.

Both RS/6000 and SP configurator sessions are required, but for the built-in safeguards to take effect, you should create simultaneous sessions to allow linkages to be formed. The SP side will provide the cables, PSSP 3.1 code, and installation instructions, as well as tracking for virtual nodes and frames, and the RS/6000 side will provide the RS/6000 SP System Attachment Adapter where attachment to the SP Switch is desired.

7.10.6.1 The SP Configurator Session

All orders for the SP attachment share several common elements: cables, ground straps, Ethernet cards for the S7A/S80, PSSP 3.1 code, AIX 4.3.3 code, and SP node and frame features that automatically trigger several of the preceding items and also provide tracking information on node, frame, and switch number usage in the SP configurator. Most orders will require the use of both SP and RS/6000 configurator sessions since the SP System Attachment Adapter is actually ordered through the RS/6000 configurator session. Substantial safeguards are built into the configurator sessions when the SP and RS/6000 configurator sessions are jointly accessed; so, it is

important to do so in a joint session if at all possible. Be sure to indicate in the RS/6000 configurator that this is an SP-attachment order. You should see a question for this if you order your S7A/S80 new. Answering in the affirmative brings in the SP rules for adapters, which are the ones in force for an SP-attach.

The SP configurator session for all variations of SP-attach includes the following:

- (# 9122) Node Attachment Feature**
An order for this feature generates an order for the 2-15 meter RS-232 cables for hardware control and S1TERM connectivity between the CWS and the SP-attached server. It also generates an order for a 10 meter ground cable. It also traps some data so that the SP configurator session can keep track of how many nodes (real and logical) you have in your system.
- (# 9310) Switch Connection Cable**
This feature is required only if the SP-attached server is switch attached. It results in the ordering of one 10 meter switch connection cable. The 10 meter cable is the only supported length at this time.
- (# 9123) Frame Attachment Feature**
This feature keeps track of how many frames are currently in your SP system. Since the S7A/S80 SP-attached server is both a logical node and a logical frame in the PSSP code logic, it is important to track this information to avoid exceeding allowable SP limits for the number of frames.
- (# 9222) Node Attachment Ethernet BNC Boot Feature**
This feature gets you a BNC cable to allow SP Ethernet communications and booting with your SP-attached server, whether switch-attached or not.
- (# 9223) Node Attachment Ethernet Twisted Pair Boot Feature**
This feature tracks that you have chosen to incorporate your SP-attached servers as part of an SP Ethernet Twisted Pair network, but it provides no twisted pair cable. As in the past, you are responsible for securing your own twisted pair Ethernet cables.
- (# 5765-D51) PSSP 3.1**
PSSP 3.1 is required for each SP-attached server license. The media features are as follows:
(# 5800) 4 mm tape
(# 5801) 8 mm tape

(# 5803) CD-ROM

The charge features are as follows:

(# 4001) PSSP on an SP

(# 4002) PSSP on an externally attached server

7.10.6.2 The RS/6000 Configurator Session

The RS/6000 configurator provides you with the ability to order your S7A/S80 with the AIX 4.3.3 operating system with the amount of I/O, storage, and memory your customer wants. There is no difference in the allowed I/O, storage, and memory capacity between the standard S7A/S80 and the SP-attached versions except for the choice of I/O and communications adapters and the placement of some adapters. You will be asked during your session if your S7A/S80 will be SP-attached. If you answer *yes*, the SP rules for adapters will replace the RS/6000 rules for adapters for the remainder of your session.

The RS/6000 configurator session is also where you will order feature # 8396, the SP System Attachment Adapter, if your intent is to have a switch-attached SP server.

7.11 I/O Adapters Supported on PCI SMP Nodes

Table 96 lists the supported communications and storage I/O adapters for the PCI-bus SMP SP nodes.

The SMP wide nodes have 10 PCI adapter slots. The SMP thin nodes have only two available PCI adapter slots. The SMP High node has 53 expansion slots.

Note

A 100BaseTX/10BaseT/10Base2 Ethernet Adapter for the SP Ethernet is integrated into the POWER3 SMP nodes and does not use a PCI slot.

Similarly, a 10BaseT/10Base2 Ethernet Adapter for the SP Ethernet is integrated into the 332 MHz SMP nodes and does not use a PCI slot.

Table 96. Supported Communications PCI Adapters on the SMP SP Nodes

Feature Code	Description
2732	Short Wave Serial HIPPI PCI Adapter
2733	Long Wave Serial HIPPI PCI Adapter

Feature Code	Description
2741	SysKonnnect SK-NET FDDI-LP SAS PCI
2742	SysKonnnect SK-NET FDDI-LP DAS PCI
2743	SysKonnnect SK-NET FDDI-UP SAS PCI
2751	S/390 ESCON Channel PCI Adapter
2920	PCI Auto LANStreamer token-ring adapter
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI Bus
2944	128-Port Asynchronous Controller, PCI Bus
2947	PCI Artic960Hx 4-Port Selectable
2962	PCI 2-port Multiprotocol Adapter
2963	Turboways 155 PCI/UTP ATM Adapter
2968	IBM 10/100 Mbps Ethernet PCI Adapter
2969	PCI Gigabit Ethernet
2985	PCI Ethernet BNC/RJ-45 Adapter
2987	PCI Ethernet AUI/RJ-45 Adapter
2988	Turboways 155 PCI MMF ATM Adapter
4959	Token-ring adapter
6310	PCI Artic960RxD Quad Digital Trunk Adapter
6311	PCI Artic960RxD Quad Digital Trunk Resource (PCI)
6204	Universal Ultra SCSI Differential PCI-bus Adapter
6205	Ultra2 SCSI Adapter
6206	PCI Single-Ended Ultra SCSI Adapter
6207	PCI Differential Ultra SCSI Adapter
6208	PCI SCSI-2 Fast/Wide Single-Ended Adapter
6209	PCI SCSI-2 Fast/Wide Differential Adapter
6215	PCI SSA Multi-Initiator/RAID EL Adapter (Withdrawn)
6230	Advanced SerialRAID Plus Adapter
6227	Gigabit Fibre Channel Adapter

7.12 Supported Levels of AIX and PSSP on SP Nodes

Table 97 on page 309 summarizes the options for running AIX and PSSP on the SP nodes.

The Control Workstation must run at least the highest level of AIX and the highest level of PSSP as any node within the SP system.

It is possible to partition an SP system that is running AIX Version 4 on the Control Workstation with PSSP 2.1 or later. There are rules and limitations on how this can be achieved and you should refer to the SP System Planning documentation for further details. It is now possible to run AIX Version 4.1, AIX Version 4.2, and AIX Version 4.3 in the same System Partition under most circumstances.

The coexistence functions in PSSP 2.2, PSSP 2.3, and PSSP 2.4 also allow versions of AIX to be mixed within the same SP system or partition. The allowable combinations are shown in Table 97.

User Space protocol across the SP Switch is only supported on SMP nodes when running PSSP 2.3 or later; otherwise, only TCP/IP traffic can be carried across the SP Switch.

Table 97. AIX and PSSP Levels Supported on SP Nodes

AIX and PSSP Level	222 MHz SMP	332 MHz SMP		POWER3-II SMP 375 MHz			POWER3 SMP 200 MHz	
	Thin	Thin	Wide	Thin	Wide	High	Thin	Wide
AIX 4.3.2								
PSSP 3.1	N	S	S	N	N	N	S	S
AIX 4.3.3								
PSSP 3.1.1	A ²	A	A	A ¹	A ¹	A	A ³	A ³

AIX and PSSP Level	222 MHz SMP	332 MHz SMP		POWER3-II SMP 375 MHz			POWER3 SMP 200 MHz	
	Thin	Thin	Wide	Thin	Wide	High	Thin	Wide
Notes: A = Available; factory and field orderable. S = Supported, but not orderable. N = Not supported. *A single 332 MHz SMP Thin Node in a drawer is supported only by AIX V4.3.2 with PSSP V3.1. ¹ AIX Version 4.3.3 with APAR IY06844, or later, with IBM Parallel System Support Programs for AIX (PSSP) Version 3.1.1 or later. ² AIX Version 4.3.3, or later, with IBM Parallel System Support Programs for AIX (PSSP) Version 3.1.1 with APAR IY03015. ³ AIX Version 4.3.2, or later, with IBM Parallel System Support Programs for AIX (PSSP) Version 3.1 with APAR IX85457.								

It is important to note that the required minimum PTF levels must be installed in each of the preceding cases. Check with the local AIX support center for this information.

7.12.1 Withdrawn Nodes and Their RS/6000 Equivalentents

Table 98 is a list of SP nodes that are not current and their stand-alone RS/6000 model equivalentents.

Table 98. SP Nodes Not Current

Node Type	RS/6000 Equivalent	Processor
Thin 1	Model 390	POWER2 (66 MHz)
Thin 2	Model 39H	POWER2 (66 MHz)
Thin P2SC	N/A	POWER2SC (120 MHz)
Thin P2SC	Model 397	POWER2SC (160 MHz)
Wide 1	Model 590	POWER2 (66 MHz)
Wide RPQ	Model 59H	POWER2 (66 MHz)
Wide 2	Model 591	POWER2 (77 MHz)
Wide P2SC	Model 595	POWER2SC (135 MHz)
High 1 (2 way)	Model R40	PowerPC 604 (112 MHz)
High 1 (4 way)	Model R40	PowerPC 604 (112 MHz)

Node Type	RS/6000 Equivalent	Processor
High 1 (6 way)	Model R40	PowerPC 604 (112 MHz)
High 1 (8 way)	Model R40	PowerPC 604 (112 MHz)
High 2 (2 way)	Model R50	PowerPC 604e (200 MHz)
High 2 (4 way)	Model R50	PowerPC 604e (200 MHz)
High 2 (6 way)	Model R50	PowerPC 604e (200 MHz)
High 2 (8 way)	Model R50	PowerPC 604e (200 MHz)

7.13 High Availability

If high availability is important, you should always consider the option of using a larger number of less powerful nodes.

When you consider the design of the SP system for a particular customer application, one of the factors usually taken into consideration is the degree of high availability required.

High availability within the SP environment is achieved in three ways:

1. By providing a base system with a high degree of reliability that is less prone to error.
2. By providing ways within the SP of achieving concurrent maintenance, where possible, to reduce the amount of time when the applications are not available due to regular maintenance and housekeeping work.
3. By providing a *failover* function to be instigated in the event of failure of any kind. This kind of concept is usually provided by using the High Availability Clustered Multi-Processing (HACMP) software, plus redundancy of components.

7.13.0.1 Redundancy

You can provide redundancy within the SP system by providing RAID or mirrored disks, alternate networks, additional adapters, and, in particular, spare or additional nodes to be used in the event of failure.

7.13.0.2 Disk Pairs - RAID

Note that there are differences between the various SP nodes with respect to their capability to provide disk and adapter redundancy for internal disks. Most nodes come with *pairs* of internal disks and allow the ROOTVG volume group to be mirrored while also allowing these internal disks to be accessed

using multiple SCSI adapters, thus removing the adapter as a potential single point of failure.

If, for example, you decide to employ one powerful wide node to run a specific application, then a failure of that node does not allow for anything else to be used. It is usually not cost-effective to provide an additional, equally performing node purely for availability purposes in these circumstances. So, this is clearly a case of “all the eggs in one basket,” which is one of the well known weaknesses of using a single SMP system to run all applications.

So if high availability is important, you can consider the option of using a number of less powerful nodes in combination so that if any single one fails, the rest will carry on regardless. It is possible to build a highly available solution in this manner without dramatically increasing the cost of the overall system.

7.13.0.3 Application Isolation

If multiple applications are run concurrently on any UNIX uniprocessor or SMP processor, because of the fact they are sharing resources, they will be in contention, and it is impossible to prevent the impact of one application on another.

For example, if applications A, B, and C are run on the same UNIX system, and application A is used very heavily, its increased workload on the system will have a detrimental effect on applications B and C. It is impossible to prevent this on a single (uniprocessor or SMP) system.

One of the benefits of the SP system is the fact that different applications can be run on different nodes within the SP system and guarantee the performance and service levels for the applications on each node separately. This is even more important, for example, when different test or development systems, or different RDBMS software, need to be run.

SAP is again a good example where this capability is critical. You can run separate application servers, database servers, update and batch servers, and even separate instances of SAP on the same SP system by utilizing multiple nodes.

So again, this is a case where you may not want to run everything on a small number of SP nodes, but may well choose to utilize a larger number of less powerful nodes to allow for application isolation.

7.14 System Partitioning

System partitioning within the SP allows the SP system to be divided into logically separate SP systems. This gives the ability to completely isolate SP environments from each other and to have, for example, two production environments or a production and a test environment.

Although SP system partitioning has been used to assist with migration from AIX Version 3 to AIX Version 4 in an SP environment in the past, this was not its primary purpose. In addition, it is now no longer supported to run AIX Version 3 on the SP system. Running combinations of AIX Version 4 on the same SP system or within a system partition is now very easy using the coexistence support in PSSP; so, you do not have to partition your SP simply for this reason alone. However, be aware that if your node is not running AIX 4.3 or later, it would be at an unsupported level of AIX.

You can implement system partitioning on any SP system that runs AIX 4.1 and PSSP 2.1 or later. This allows us to divide up the SP system into separate logical SP systems so that you can run with completely separate environments within the same SP. In particular, apart from nodes being isolated from each other if they are in separate partitions, you also have separate logical SP Switches within the SP system, and SP Switch traffic in one partition does not contend with other partitions.

At the time of publication, system partitioning is less used, and there are only a few environments where the function is required.

The partitioning is achieved according to certain defined rules, which are based on Switch Chip Boundaries. The SP Switch adapters in each node in a frame are physically cabled into the SP Switch at the bottom of the SP frame and are connected to various Switch Chips according to their position in the frame.

7.15 Peripheral Devices

The attachment of peripheral devices, such as disk subsystems, tape drives, CD-ROMs, and printers, is very straightforward on the SP. There are no SP-specific peripherals; since the SP uses mainstream RS/6000 node technology, it simply inherits the array of peripheral devices available to the RS/6000 family. The SPs shared-nothing architecture gives rise to two key concepts when attaching peripherals:

1. Each node has I/O slots. Think of each node as a stand-alone machine when attaching peripherals: It can attach virtually any peripheral available

to the RS/6000 family, such as SCSI and SSA disk subsystems, MagStar tape drives, and so on. Peripheral device attachment is very flexible because each node can have its own mix of peripherals, or none at all.

2. From an overall system viewpoint, as nodes are added, I/O slots are added, thus the scalability of I/O device attachment is tremendous. A 512-node high node system would have several thousand I/O slots!

There are some general considerations for peripheral device attachment:

- Since the nodes are housed in frames, cable lengths from the I/O adapter in a node slot to the actual peripheral must be long enough.
- Devices, such as CD-ROMs and bootable tape drives, may be attached directly to SP nodes, but IBM does not support their use as installation devices. Nodes must be network-installed by the CWS or a boot/install server.
- Many node types do not have serial ports. High nodes have two free serial ports for general use.
- Graphics accelerators for attachment of displays are not supported.

Chapter 8. IBM NUMA-Q

The NUMA-Q brand of IBMs Enterprise Server Group, formerly Sequent Computer Systems, is the leader in scalable Intel-based server solutions for e-business. IBM NUMA-Q platforms and services are optimized for the scalability, availability, and manageability requirements of large and rapidly growing e-business infrastructures. There are more than 10,000 IBM NUMA-Q installations worldwide, including some of the world's largest and most sophisticated business intelligence, customer relationship management (CRM), and enterprise resource planning (ERP) environments.

In this chapter, the IBM NUMA-Q server technology and products are discussed.

8.1 NUMA Introduction

The Non-Uniform Memory Access (NUMA) computer architecture has proven to be the logical growth path for SMP in the e-business era. NUMA was born as a Stanford University research project in the late 1980s and was first offered as a commercial server in 1996 by Sequent Computer Systems. Today, NUMA systems run some of the largest UNIX database applications in the world and are gaining acceptance as a mainstream technology for e-business.

NUMA offers significant benefits to fast emerging e-business environments, including massive scalability of processing power and I/O, high availability, broad flexibility for work load and resource management, and no change for the SMP programming model.

NUMA achieves its scalability and flexibility benefits by linking multiple small SMP hardware building blocks containing processors and memory with a fast interconnect in a single system image. The local memory on each building block is used to cache frequently accessed data to limit the frequency of non-local memory accesses. Current NUMA architectures can support hundreds of processors in a single system using this modular approach.

8.1.1 IBM NUMA-Q Architecture

IBM NUMA-Q systems use the cc-NUMA or cache coherent approach to NUMA. In a cc-NUMA system, memory is located close to a number of processors, for speed, but still physically tied together to form a single memory image. Cache coherent refers to the requirement for keeping multiple copies of data up to date or for transferring data between multiple instances

of the OS or application. Hardware cache coherent means that there is no software requirement for keeping multiple copies of data up to date or for transferring data between multiple instances of the OS or application. In IBM's NUMA-Q system, this work is all managed at the hardware level, just as in any SMP node, with a single instance of the OS and multiple processors.

NUMA-Q is a collection of four Intel-based processor building blocks (named quads) interconnected through a set of high-speed point-to-point links (IQ-Link). Each building block is a complete SMP machine containing processors, cache, memory, and I/O busses, which are uniquely arranged. Units of memory are located close to the processors to maximize overall system performance. I/O is also near each processor so that all memory and I/O accesses can be satisfied within the quad. Multiple quads are interconnected with a hardware-based, cache-coherent, very high speed Scalable Coherent Interface (SCI) connection to create single NUMA systems, in the same way that processor boards are added to a backplane of a conventional SMP system.

8.1.2 IQ-Link

IQ-Link, IBM's innovative technology for linking multiple quads, enables data transfer between quads at speeds up to 1 GB/s. The IQ-Link is not a backplane; however, it is a daisy chain connection between quads. IQ-Link enables high performance by cutting latencies to those of the previous generation and providing cache coherency within each system node. Flexible memory configurations of 1 GB to 64 GB per node provide optimal memory sizing to meet performance requirements. Because IQ-Link can provide a unified view of memory, with ranges of address space parceled out to quads, one instance of the OS and the applications may simultaneously run on all interconnected quads.

8.1.3 NUMA-Q SAN

NUMA-Q is the first UNIX-based system to offer integral MVS-style multipathing switched Fibre Channel SAN (Storage Area Network). This industry-leading capability is a key enabling technology for e-business and CRM applications requiring support for very large, high performance transaction and data warehouse environments. The Fibre Channel SAN allows large backend UNIX database machines and hundreds of front end UNIX or NT application servers to use a common switched fabric and cost effectively share data center-class disk storage and tape libraries.

The NUMA-Q runs I/O directly to its connected storage devices over a switched Fibre Channel SAN fabric, and not over the interconnect that

handles memory accesses. On the NUMA-Q systems, this eliminates the resource contention that reduces throughput in large SMP systems as processors are added. Also, because of the I/O multi-pathing that is supported at the operating system level, NUMA-Q offers inherently fault tolerant SAN.

8.1.4 NUMACenter

The NUMACenter is a Windows NT application deployment framework for e-business supporting mixed Windows NT and UNIX server and Storage Area Network (SAN) resources in a single environment.

IBMs NUMACenter platform allows administrators to have full control of the migration of NT into their data center by allowing them to run concurrent UNIX and NT front-end and back office applications. NUMACenter's all NT data center provides NT customers with an environment that is pre-integrated and characterized across enterprise-level peripheral and software stacks.

8.1.5 A Defining Technology

On July 1999, IBM announced a definitive merger agreement with Sequent calling NUMA "a defining technology for early 21st century UNIX and NT servers". As the pressures of e-business weigh on IT operations, larger IT vendors have begun to develop and announce plans for their own approaches to the NUMA architecture.

8.2 NUMA-Q Products, Solutions, and Services

IBM NUMA-Q designs, manufactures, sells, and supports Intel-based data center systems and solutions as described in this section.

8.2.1 NUMA-Q 2000 Enterprise Server

The NUMA-Q 2000 is a highly scalable enterprise server designed to deliver data center class performance and high availability. NUMA-Q 2000 utilizes Intel Pentium III Xeon processors, enabling increased performance and scalability. The NUMA-Q architecture protects infrastructure investments, allowing successive generations of Intel processors to interact efficiently in the same system. NUMA-Q 2000 is currently optimized for up to 64 processors. The architecture supports up to 256 processors.

8.2.2 DYNIX/ptx

DYNIX/ptx is among the most robust and reliable implementations of UNIX-like OS for commercial SMP and NUMA systems running

enterprise-level applications. DYNIX/ptx is designed for the Intel Architecture environment, and provides both a data-center-proven 32-bit UNIX-based OS and a low risk path to 64-bit computing. Many of the distinguishing technologies of DYNIX/ptx are being migrated into IBMs 64-bit Project Monterey UNIX implementation for Intel's IA-64 Itanium processor.

8.2.3 Specifications of NUMA-Q 2000 Enterprise Server

The following details are provided for a brief overview on the IBM NUMA-Q server technology and products. For a detailed proposal template, presentations, and configuration tool, visit the Web site of IBM NUMA-Q Knowledge Repository called Sequent Coporate Electronic Library (SCEL). There is a form on the IBM Server Sales Site that can be used to request a SCEL account. Go to:

<http://scel.sequent.com>

Click the help link at the bottom of the provided page. Follow the directions on that page to find the form. After submitting the form, you will receive an e-mail confirmation letter with your password and further directions.

The specifications of the NUMA-Q 2000 Quad Model 410 are provided in Table 99 and the specification for the cabinet summary are provided in Table 100.

Table 99. Specification Summary of IBM NUMA-Q 2000 Model E410

NUMA-Q 2000 Node Model 410	
Number of Processors	Single Node: 4 to 64 Two Node Cluster: 8 to 128
Architecture	Intel Pentium III Xeon processor, 700 MHz, 1 MB, 2 MB
Operating System	DYNIX/ptx
Processor Cache	L1: 32 KB; L2: 1 MB, 2 MB
Memory Options	
Single Node	1 to 64 GB
Two Node Cluster	2 to 128 GB
Storage Options	
Internal Storage	32X CD-ROM, maximum 1 x 36 GB disk, QIC525 Tape Drive

NUMA-Q 2000 Node Model 410	
External Storage Options	EMC: maximum 36 TB 12 each 3630s or 3830s or 6 each 3930s, CLARiiON: maximum 25 TB 1440 drives, IBM Enterprise Storage Server
I/O Subsystem	PCI (Peripheral Component Interconnect)
PCIs Supported	SCSI-2, Fibre Channel, Ethernet, ATM, FDDI, CDDI, token ring, Synchronous Controller
PCI Channels/Buses	2 at 266 MBps (64-bit PCI)
Total PCI Slots per Quad	7
I/O Bandwidth	Single Node: maximum 320 MBps
	Two Node Cluster: maximum 5120 MBps

Table 100. IBM NUMA-Q Cabinet

IBM NUMA-Q Cabinet	
Workgroup Cabinet	Enterprise Cabinet
200 VAC single-phase, 16A	220 VAC 3-phase, 30A
3.1 kVA max consumption	5.3 kVA max consumption
10,000 BTU/hour max heat dissipation	17,200 BTU/hour max heat dissipation
AC distribution strip (20A outlets)	Phase-balanced AC distribution (20A outlets)
Maximum weight: 1000 lbs	Maximum weight: 1600 lbs
Cabinet Dimensions	
Workgroup Cabinet	Enterprise Cabinet
Height: 42.325"	Height: 70.75"
Width: 29.30"	Width: 29.30"
Depth: 46.328"	Depth: 46.25"

Table 101 summarizes the IBM NUMA-Q 2000 Enterprise Server, which is expandable from 1 to 64 quads.

Table 101. NUMA-Q Server Summary

NUMA-Q Server	NQ2000-01	NQ2000-2L	NQ2000-02	NQ2000-03	NQ2000-04	NQ2000-16
Number of Processors - Latest Pentium III Xeon	1 Quad System	2 Quad System	2 Quad System	3 Quad System	4 Quad System	16 Quad System
Cache: Level 1; Level 2	32 KB; 1 MB or 2 MB	32 KB; 1 MB or 2 MB	32 KB; 1 MB or 2 MB	32 KB; 1 MB or 2 MB	32 KB; 1 MB or 2 MB	32 KB; 1 MB or 2 MB
Main Memory	Up to 8 GB	Up to 16 GB	Up to 16 GB	Up to 24 GB	Up to 32 GB	Up to 64 GB
Quad Interconnect	N/A	Lower Cost Direct Link 2 Card	Quad-to-Quad direct connect supported or IQ Ring Programmable Control Switch	IQ Ring Programmable Control Switch	IQ Ring Programmable Control Switch	IQ Ring Programmable Control Switch
IQ Link Cache Options	N/A	N/A	128 MB Cache Expander*	128 MB Cache Expander*	128 MB Cache Expander	128 MB Cache Expander
Storage Options						
Internal Storage (Boot, SW Loading,?)	18 GB or 36 GB Disk, 32 X CD-ROM, QIC525 Tape					

NUMA-Q Server	NQ2000-01	NQ2000-2L	NQ2000-02	NQ2000-03	NQ2000-04	NQ2000-16
Max. Storage Cap.	552 TB	552 TB	552 TB	552 TB	552 TB	552 TB
I/O Subsystem						
I/O Interconnect	Peripheral Component Interconnect	Same as 1 Quad	Same as 1 Quad	Same as 1 Quad	Same as 1 Quad	Same as 1 Quad
PCI Channels/Buses	2 @ 266 MBps (64-bit PCI, 32-bit supported)	4 @ 266 MBps (64-bit PCI, 32-bit supported)	4 @ 266 MBps (64-bit PCI, 32-bit supported)	6 @ 266 MBps (64-bit PCI, 32-bit supported)	8 @ 266 MBps (64-bit PCI, 32-bit supported)	32 @ 266 MBps (64-bit PCI, 32-bit supported)
I/O Bandwidth (1)	200 MBps	400 MBps	400 MBps	600 MBps	800 MBps	3200 MBps
Number of PCI slots	7	14	14	21	28	112
Peripheral Choices and Connectivity Options						
IBM ESS (Shark)	IBM ESS Storage can now be attached to all NUMA-Q Servers (2)					
EMC 3630	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge
EMC 3830	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge
EMC 3930	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge, Direct Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge	SAN Only, SAN Bridge
CLARiiON DAE	Direct FC	Direct FC	Direct FC			
HDS 5800	SAN Only, Direct FC	SAN Only, Direct FC	SAN Only, Direct FC	SAN Only	SAN Only	SAN Only
IBM 3590 Tape Drive	IBM 3590 Tape Drive can now be attached to all NUMA-Q Servers (2)					
IBM 3494 Tape Library	IBM 3494 Tape Library can now be attached to all NUMA-Q Servers (2)					
DLT 7000 Table-top	Direct SCSI, SAN Bridge, Direct Bridge	Direct SCSI, SAN Bridge, Direct Bridge	Direct SCSI, SAN Bridge, Direct Bridge	SAN Bridge	SAN Bridge	SAN Bridge
STK 9730	Direct SCSI, SAN Bridge, Direct Bridge	Direct SCSI, SAN Bridge, Direct Bridge	Direct SCSI, SAN Bridge, Direct Bridge	SAN Bridge	SAN Bridge	SAN Bridge

NUMA-Q Server	NQ2000-01	NQ2000-2L	NQ2000-02	NQ2000-03	NQ2000-04	NQ2000-16
STK 9710/9740	SAN Bridge, Direct Bridge	SAN Bridge, Direct Bridge	SAN Bridge, Direct Bridge	SAN Bridge	SAN Bridge	SAN Bridge
Applicable to all NUMA systems						
*128 MB Cache Expander is required when memory per Quad exceeds 4 GB.						
(1) 200 MBps I/O bandwidth assumes NQ Quad is limiting factor (with four PCI FC cards); two PCI FC cards in a Quad provide 160 MBps						
(2) NUMA-Q interface to IBM Disk and Tape requires special software drivers and cables - refer to NUMA-Q configuration tool						
SAN = Storage Area Network; SAN Only = pure Fibre Channel (FC) connection with FC Switch						
SAN Bridge = FC Switch plus FC Bridge to SCSI; Direct FC = PCI FC connect direct to peripheral with no FC Switch						

Chapter 9. Communications and Storage I/O Adapters

This chapter describes the communication and storage I/O adapters that are currently available with the RS/6000 systems. Table 102 lists the available PCI I/O adapters.

Table 102. Available PCI I/O Adapters and their Feature Codes

Feature Code	PCI Adapter Name
2493	SCSI-2 F/W RAID
2494	PCI 3-Channel Ultra2 SCSI RAID Adapter
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T
2732	IBM Short-wave Serial HIPPI PCI Adapter for RS/6000
2733	IBM Long-wave Serial HIPPI PCI Adapter for RS/6000
2741	SysKonnnect SK-NET FDDI-LP SAS PCI for PowerPC Systems
2742	FDDI SK-NET LP DAS
2743	FDDI SK-NET UP SAS
2751	S/390 ESCON Channel Adapter
2943	8-Port Asynchronous Adapter EIA 232/RS-422
2944	128-Port Asynchronous Adapter EIA-232
2947	ARTIC960Hx 4-Port Selectable Adapter
2948	ARTIC960Hx 4-Port T1/E1 PCI
2962	2-port Multiprotocol X.25 Adapter
2963	ATM 155 Turboways UTP Adapter
2968	Ethernet 10/100 MB
2969	Gigabit Ethernet - SX
2985	PCI Ethernet BNC/RJ-45 Adapter
2987	PCI Ethernet AUI/RJ-45 Adapter
2988	ATM 155 Turboways MMF Adapter
2998	ATM 25 Turboways
4951	4-Port 10/100 Base-TX Ethernet

Feature Code	PCI Adapter Name
4959	4/16 Mbps token-ring adapter
6204	PCI Universal Differential Ultra SCSI
6205	PCI Dual Channel Ultra2 SCSI Adapter
6206	Ultra SCSI Single-Ended
6207	Ultra SCSI Differential
6208	SCSI-2 F/W Single-Ended
6209	SCSI-2 F/W Differential
6215	SSA RAID 5 (accepts optional SSA Fast-Write Cache module (# 6222))
6227	Gigabit Fibre Channel Adapter
6230	Advanced SerialRAID Plus Adapter
6310	ARTIC960RxD Quad Digital Trunk Adapter
6311	IBM ARTIC960RxF PCI Adapter for RS/6000
8396	IBM RS/6000 SP System Attachment Adapter

The PCI adapters are categorized into following groups:

- ISDN adapter
- System adapters
- FDDI adapters
- Asynchronous adapters
- ARTIC adapters
- WAN adapters
- ATM adapters
- Ethernet and token-ring adapters
- PCI Storage adapters

9.1 ISDN Adapter

The following is a description of the available ISDN adapter.

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

The Eicon ISDN DIVA PRO 2.0 S/T Adapter (# 2708) 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 (# 2708) provides basic rate 2B + D with 64 Kbps on each B channel. It requires a connection to Network Terminator -1 (NT-1) from the phone line provider. TCP/IP applications are supported.

Note

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

9.2 System Adapters

The following is a description of the available system adapters.

9.2.1 IBM Short-wave Serial HIPPI PCI Adapter (# 2732)

The IBM Serial HIPPI PCI Adapter is a single slot, full size, 32-bit PCI adapter supporting TCP/IP over a serial HIPPI link. Data is sent and received over optical fiber at 1.2 Gbps using the HIPPI standard 20/24-bit encoding scheme with short wave optical transducers. The effective maximum data rate of the HIPPI interface is 800 Mbps.

Serial HIPPI PCI Connectivity Version 4.3 AIX LPP (5765-E07) with the IBM Serial HIPPI PCI Adapter (# 2732) interconnects the RS/6000 and serial HIPPI networks. This solution supports connection to serial HIPPI networks using TCP/IP.

The IBM HIPPI Protocol Services Program (5765-E40) provides additional protocols for high-speed communications and storage to an RS/6000 system using a PCI channel (bus) and is supported on AIX Versions 4.3.2 and 4.3.3.

The protocols, IPI-3, NDA, and FP, are supported using the IBM HIPPI Protocol Services program, with the IBM Serial HIPPI PCI Connectivity Version 4.3 program and the IBM Serial HIPPI PCI Adapter.

9.2.1.1 Software Requirements

- AIX 4.3.2, or later
- IBM HIPPI PCI Connectivity Version 4.3 AIX LPP (5765-E07)

9.2.2 IBM Long-wave Serial HIPPI PCI Adapter (# 2733)

The IBM Serial HIPPI PCI Adapter is a single slot, full size, 32-bit PCI adapter supporting TCP/IP over a Serial HIPPI link. Data is sent and received over optical fiber at 1.2 Gbps using the HIPPI standard 20/24 bit encoding scheme with long wave optical transducers. The effective maximum data rate of the HIPPI interface is 800 Mbps.

Serial HIPPI PCI Connectivity Version 4.3 AIX LPP (5765-E07) with the IBM Serial HIPPI PCI Adapter (# 2733) interconnects the RS/6000 and serial HIPPI networks. This solution supports connection to serial HIPPI networks using TCP/IP.

The IBM HIPPI Protocol Services Program (5765-E40) provides additional protocols for high-speed communications and storage to an RS/6000 system using a PCI channel (bus) and is supported on AIX Versions 4.3.2 and 4.3.3.

The protocols, IPI-3, NDA, and FP, are supported using the IBM HIPPI Protocol Services program, with the IBM Serial HIPPI PCI Connectivity Version 4.3 program and the IBM Serial HIPPI PCI Adapter.

9.2.2.1 Software Requirements

- AIX 4.3.2, or later
- IBM HIPPI PCI Connectivity Version 4.3 AIX LPP (5765-E07)

9.2.3 S/390 ESCON Channel Adapter (# 2751)

The PCI S/390 ESCON Channel Adapter (# 2751) provides the ability to attach to IBM Enterprise Systems Connection (ESCON) channels on System/390 mainframes or to attach IBM ESCON tape devices through channel emulation. This direct ESCON channel connection provides a fiber optic link that can take advantage of ESCON Directors (fiber optic switches) permitting multiple channel connections.

9.2.3.1 Feature Characteristics

- Full length PCI adapter
- PCI 32-bit Bus Master adapter
- Mainframe connectivity features include:

- Supports attachment to either 10 MB or 17 MB ESCON channels
- Supports VM, MVS, and OS/390
- Supports CLIO/S
- Supports ESCON Multiple Image Facility (EMIF)
- Maximum distance supported, 43 km using LED and XDF ESCON links
- S/390 TCP/IP for VM and MVS
- Tape subsystem attachment features include:
 - Supports 3490 and 3490E ESCON attachable tape subsystems
 - Supports 3590-B01, B11, E01, and E11 tape drives; and 3590-A00, A50, and A60 control units (3590 control unit required)
 - Supports 3494 and 3595 Tape Library Data Servers
 - Maximum distance supported, 23 km using two ESCON Directors with XDF

9.2.3.2 Feature Components

- Installation Instructions
- Diagnostic wrap plug (diagnostics are included in AIX Diagnostic CD-ROM)
- CD-ROM with device drivers shipped with the LPPs
- Instruction manual shipped with the LPP

9.2.3.3 Customer-Supplied Components

- ESCON cabling, requires 62.5/125 multimode fiber cable with ESCON duplex connectors on both ends
- One of the following:
 - AIX program feature, PCI ESCON Control Unit Connectivity V2 (5765-D49) for mainframe connectivity
 - AIX program feature, PCI ESCON Tape Attachment V2 (5765-E04) for tape subsystem attachment

9.2.3.4 Software Requirements

- Supported with AIX 4.2.1 and above on selected hardware platforms
- One of the following:
 - PCI ESCON Control Unit Connectivity V2 (separately ordered as LPP 5765-D49)

- PCI ESCON Tape Attachment V2 (separately ordered as LPP 5765-D49) for tape subsystem attachment

9.2.3.5 Floor Raceways

Raised floor installations require under-floor raceways. For non-raised floor installations, IBM recommends raceways to be installed to protect cables from being damaged.

9.2.4 IBM RS/6000 SP System Attachment Adapter (# 8396)

The S-series servers can function as an attached SMP server within the IBM RS/6000 SP environment operating under the control of the IBM Parallel Systems Support Programs (PSSP) for AIX. The joining of these technologies satisfies the need many SP customers have for large, powerful, and memory-rich processors for their database servers and ERP applications and generally provides a single point-of-control for the entire system. This interconnection can be accomplished utilizing the IBM RS/6000 SP System Attachment Adapter (switch-attachment) or with an Ethernet connection (switchless connection).

The RS/6000 SP System Attachment Adapter (#8396) must always be located in slot 10 of the primary I/O drawer. No adapters may be installed in slot 9 or 11 of the primary drawer when the SP attachment adapter is installed. Only one SP System Attachment Adapter (# 8396) is allowed per S-series system.

9.2.4.1 Hardware Requirements

- One free PCI slot
- RS/6000 SP system
- Interconnection cables
- One 10 Mbps Ethernet adapter
- IBM Parallel System Support Program for AIX Version 3.1, or later

9.2.4.2 Software Requirements

- AIX 4.3
- PSSP V3.1, or later

Note

Refer to the *Implementing the SP-attached, S70 Advanced or S80 Server* white paper or contact your local IBM representative for additional planning information.

9.3 FDDI Adapters

The following sections discuss the available FDDI adapters.

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

The SysKonnnect SK-NET FDDI-LP SAS PCI Adapter (# 2741) is a fiber optical FDDI Single Attach Station that is compatible with the FDDI-ANSI X3T12 specifications and FDDI Standard Series. The adapter provides single attachment to an FDDI concentrator (or point-to-point) using fiber optic cabling (not supplied with the adapter).

9.3.1.1 Feature Characteristics

- Supports single-ring FDDI attachment at 100 Mbps using a customer-supplied FDDI concentrator
- Supports all TCP/IP protocols and ANSI Station Management 7.3

9.3.1.2 Feature Components

- Adapter card
- Diagnostic wrap plug
- Diskette with adapter device driver
- Installation instructions

The SysKonnnect SK-NET FDDI-LP SAS PCI Adapter supports Remote IPL and NIM with the appropriate system firmware. A configuration using more than one SysKonnnect FDDI adapter (# 2741, # 2742, or # 2743), where any one of them is # 2743 (SysKonnnect SK-NET FDDI-UP SAS PCI), constitutes a Class A system.

9.3.1.3 Customer-supplied Components

- A FDDI concentrator, such as the IBM 8240 (or equivalent) concentrator, to connect to your FDDI local area network
- One 62.5/125 micron multimode fiber duplex cable with SC connectors

9.3.1.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver and FDDI common code (provided with adapter)

Note

Refer to the *PCI Adapter Placement Reference*, SA38-0538, for important information regarding permissible and optimum slot placement of your PCI adapters.

9.3.2 FDDI SK-NET LP DAS Adapter (# 2742)

The SysKonnnect SK-NET FDDI-LP DAS PCI Adapter (# 2742) is a fiber optical FDDI Dual Attach Station that is compatible with the FDDI-ANSI X3T12 specifications and FDDI Standard Series. The 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).

9.3.2.1 Feature Characteristics

- Supports dual ring FDDI attachment at 100 Mbps
- Supports all TCP/IP protocols and ANSI Station Management (SMT) 7.3

9.3.2.2 Feature Components

- Adapter card
- Diagnostic wrap plug
- Diskette with adapter device driver
- Installation instructions

The SysKonnnect SK-NET FDDI-LP DAS PCI Adapter supports Remote IPL and NIM with the appropriate system firmware. A configuration using more than one SysKonnnect FDDI adapter (# 2741, # 2742, or # 2743) where any one of them is # 2743 (SysKonnnect SK-NET FDDI-UP SAS PCI), constitutes a Class A system.

9.3.2.3 Customer-Supplied Components

- A FDDI concentrator, such as the IBM 8240 (or equivalent) concentrator, to connect to the FDDI network for dual homing configurations
- Two 62.5/125 micron multimode fiber duplex cables with SC connectors

9.3.2.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver and FDDI common code (provided with adapter)

Note

Refer to the *PCI Adapter Placement Reference*, SA38-0538, for important information regarding permissible and optimum slot placement of your PCI adapters.

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

The SysKonnnect SK-NET FDDI-UP SAS PCI Adapter (# 2743) is a Copper FDDI Single Attach Station that is compatible with the FDDI-ANSI X3T12 specifications and FDDI Standard Series. The adapter provides single attachment to an FDDI concentrator (or point-to-point) using Category 5 Unshielded Twisted Pair cabling (not supplied with the adapter).

9.3.3.1 Feature Characteristics

- Supports single ring FDDI attachment at 100 Mbps
- Supports all TCP/IP protocols and ANSI Station Management (SMT) 7.3

9.3.3.2 Feature Components

- Adapter card
- Diagnostic wrap plug
- Diskette with adapter device driver
- Installation instructions

The SysKonnnect SK-NET FDDI-UP SAS PCI Adapter supports remote IPL and NIM with the appropriate system firmware. A configuration using more than one SysKonnnect FDDI adapter (# 2741, # 2742, or # 2743), where any one of them is # 2743 (SysKonnnect SK-NET FDDI-UP SAS PCI), constitutes a Class A system.

9.3.3.3 Customer-Supplied Components

- A FDDI concentrator, such as the IBM 8240 (or equivalent) concentrator, to connect to the FDDI network for dual homing configurations
- One Unshielded Twisted Pair Category 5 cable

9.3.3.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver and FDDI common code (provided with adapter)

Note

Refer to the *PCI Adapter Placement Reference*, SA38-0538, for important information regarding permissible and optimum slot placement of your PCI adapters.

9.4 Asynchronous Adapters

The following sections discuss the available asynchronous adapters.

9.4.1 8-Port Asynchronous Adapter EIA 232/RS-422 (# 2943)

The 8-port Asynchronous Adapter (# 2943) provides up to eight EIA-232 or RS-422 asynchronous serial lines from a single PCI bus slot. This adapter adheres to the Peripheral Component Interconnect (PCI) Revision 2.1 standards for EIA-232 and RS-422. It features a low cost, high-performance 32-bit card, 33 MHz bus speed, and a PCI bus transfer rate of 132 MBps.

This adapter provides a single DB-78 output which connects directly to the 8-port DB-25 connector box. All eight ports are software programmable to support both EIA-232 and RS-422 protocols at baud rates up to 230 Kbps. The full set of modem control lines for asynchronous communication are provided for each port. Devices such as terminals, modems, processors, printers, and controllers may be attached.

9.4.1.1 Feature Characteristics

- 8-port asynchronous device connections
- 32-bit Bus Master PCI bus (132 MBps)
- Short-form factor PCI Adapter
- EIA-232 maximum distance 31 m and 62 m dependent on baud rate and RAN
- RS-422 maximum distance 1200 m dependent on baud rate
- 230 Kbps maximum baud rate
- Supports TxD, RxD, RTS, CTS, DSR, DCD, DTR, and RI on EIA 232

- Supports +TxD, -TxD, +RxD, and -RxD on RS-422

9.4.1.2 Feature Components

- Adapter card
- 25-pin diagnostic wrap plug
- Diskette with adapter device driver
- Installation instructions
- Includes external 3 m DB78 cable to 8-port DB25 breakout box

9.4.1.3 Customer-Supplied Components

A 3 m cable with attached breakout box is supplied with each adapter. The customer must supply all cables needed to connect peripheral equipment to this adapter.

9.4.1.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver LPP image (provided with adapter)

9.4.2 128-Port Asynchronous Adapter EIA-232 (# 2944)

The 128-port asynchronous adapter (# 2944) provides up to 128 EIA-232 asynchronous serial lines from a single PCI bus slot. This adapter adheres to the Peripheral Component Interconnect PCI standard. It features a low cost, high performance 32-bit card, 33 MHz bus speed, and a PCI bus transfer rate of 132 MBps.

Two 2.4 Mbps synchronous channels link the adapter to a maximum of eight 16-port Remote Asynchronous Nodes (RANs). Each synchronous channel uses an HD-15 female connector to link up to four RANs. Each RAN supports EIA-232 or RS-422 connections (16 per RAN) and up to eight RANs may be connected together yielding a total of 128 ports. If you use a 232 RAN, then all the ports on that RAN only support this standard. It is the same for RS-422 RANs. The RAN defines the granularity that you have in choosing which standard. The RAN utilizes an RJ-45 connector to provide interface signals at speeds up to 230 Kbps at a limited number of ports.

9.4.2.1 Feature Characteristics

- 32-bit Bus Master PCI bus
- Two synchronous channels to RAN

- EIA-232 maximum distance 31 m and 62 m dependent on baud rate and RAN
- RS-422 maximum distance 1200 m dependent on baud rate

9.4.2.2 Customer-Supplied Components

Feature (# 2944) utilizes the following optional Remote Asynchronous Nodes (RANs) and device cables, which are available from IBM. Remember when upgrading from the older 64-port Asynchronous Adapter, the RANs will have to be changed, and RJ-45 to DB-25 converter cables will be required.

Table 103 and Table 104 provide the related feature codes and their descriptions.

Table 103. The 1.2 Mbps RANs and Cables

Feature Code	Description
8131	128-port Asynchronous Controller Node Cable, 4.5 m
8132	128-port Asynchronous Controller Cable 23 cm (9 in.) ¹
8133	RJ-45 to DB-25 Converter Cable
8136	Rack Mountable Remote Asynchronous Node 16-Port EIA-232
Note: ¹ This cable may be substituted for the 4.5 meter (15 foot) async controller cable whenever a customer configuration requires stacked Remote Async Nodes.	

Table 104. The 2.4 Mbps RANs and Cables

Feature Code	Description
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232
8138	Enhanced Remote Asynchronous Node 16-Port RS-422
2934	Asynchronous Terminal/Printer Cable, EIA-232
3926	Async Printer/Terminal Cable, 9-pin to 25-pin, 4 M
3124	Serial to Serial Port Cable for Drawer/Drawer
3125	Serial to Serial Port Cable for Rack/Rack

9.4.2.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP Nodes

- Adapter device driver LPP image (provided with adapter)

9.5 ARTIC Adapters

The following sections discuss the available ARTIC adapters.

9.5.1 ARTIC960Hx 4-Port Selectable Adapter (# 2947)

The ARTIC960Hx 4-Port Selectable PCI Adapter (# 2947) is a one-slot, standard-length, 32-bit PCI card. It provides four ports supporting either EIA-232, EIA530, RS-449, X.21, or V.35 specifications. Only one standard can be used at a time. Each port supports speeds up to 2.0 Mbps. The adapter provides hardware to support WAN protocols, such as X.25, SNA, and Bisync.

Software support is provided by ARTIC960 Support for AIX, Developer's Kit, AIX Versions 4.2.1, 4.3.2, or later (for SDLC or Bisync support), and AIX Version 4.1.5, 4.2.1, 4.3.1, or later (for AIXLink X.25 LPP Version 1.1.5 support). The adapter can also be used for real-time device control, telephony signaling, and custom serial communication protocols.

This adapter is also equipped with a high-performance, eight-channel DMA controller. This DMA controller supports intelligent DMA operations, such as data buffer chaining and end-of-frame processing, to support high-performance communications protocols and high-throughput applications. The DMA controller is fully programmable for OEM and third-party device drivers.

9.5.1.1 Feature Characteristics

- One 120-pin port
- Supports up to four connections of the same type
- Data transfer rates of up to 2 Mbps
- Supported interfaces are:
 - EIA-232
 - EIA-530
 - RS-449
 - X.21
 - V.35
- Support for SDLC and X.25 full-duplex, synchronous protocols

9.5.1.2 Customer-Supplied Cables

A connecting cable is required. The cables listed in Table 105 are available from IBM.

Table 105. ARTIC960Hx Cables

Feature Code	Cable Description
2861	ARTIC960Hx 4-port EIA-232 cable
2862	ARTIC960Hx 4-port RS-449 cable
2863	ARTIC960Hx 4-port X.21 cable
2864	ARTIC960Hx 4-port V.35 (DTE) cable
2865	ARTIC960Hx 4-port EIA-530 cable

9.5.1.3 Software Requirements

- AIX 4.2.1 and APAR IX81861, or later
- AIX 4.3.2 and APAR IX81860, or later, for SDLC or Bisync
- AIX 4.3.3
- Adapter device driver (provided with adapter)

The LPP toolkit and IBM ARTIC960 Support Software Version 1.2 for AIX will be on a CD and will be shipped with the adapter. They will also be available on the IBM ARTIC Web site at:

<http://WWPRODSOLN.bocaron.ibm.com/artic/>

Note

The combination of features # 2947, # 2948, and # 6310 must not exceed eight.

9.5.2 IBM ARTIC960Hx 4-Port T1/E1 Adapter (# 2948)

The IBM ARTIC960Hx 4-Port T1/E1 PCI Adapter (# 2948) is a one-slot, standard-length, 32-bit PCI card. The adapter provides four ports of channelized T1/E1. This adapter also provides SCbus external TDM interface and switch and is intended for use with the Direct Talk for AIX Version 2.1 LPP application software.

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

The 4-port T1/E1 mezzanine card supports four ports running simultaneously at a maximum data rate of 2.048 Mbps duplex. Table 106 presents the maximum port speed supported for each electrical interface.

Table 106. Maximum Port Speed for Each Electrical Interface

Electrical Interface	Maximum Speed (per port)
T1	1.544 Mbps
E1	2.048 Mbps

Note

The combination of features # 2947, # 2948, and # 6310 must not exceed eight.

9.5.3 ARTIC960RxD Quad Digital Trunk Adapter (# 6310)

The ARTIC960RxD Quad Digital Trunk Adapter (# 6310) provides voice processing for up to four T1 or E1 digital trunk lines, providing connectivity for 96 (T1) or 120 (E1) voice channels in a single PCI slot. The voice processing function is provided by DirectTalk for AIX, Version 2.1 LPP. The adapter provides high-function control of I/O operations and serves to off-load I/O tasks from the system microprocessor.

9.5.3.1 Feature Characteristics

- 32-bit PCI 2.1 adapter with one 36-pin, high-density port
- Support for up to four T1 or E1 trunk lines
- Supports voice processing using DirectTalk for AIX

9.5.3.2 Feature Components

- One ARTIC960RxD Adapter (# 6310)
- A connecting cable (as listed in Table 107) is required.

Table 107. ARTIC960 Cables and Features

Feature Code	Description
2709	ARTIC960Hx 4-port T1 RJ45 cable
2710	ARTIC960Hx 4-port E1 RJ45 cable
2871	ARTIC960RxD Quad DTA, T1, 100 ohm, 3 m 4-port cable
2872	ARTIC960RxD Quad DTA, T1, 100 ohm, 15 m extension cable

Feature Code	Description
2873	ARTIC960RxD Quad DTA, E1, 120 ohm balanced, 3 m 4-port cable
2874	ARTIC960RxD Quad DTA, E1, 120 ohm balanced, 7.5 m extension cable
2875	ARTIC960RxD Quad DTA, E1, 75 ohm unbalanced-grounded, 1.8 m 4-port cable
2876	ARTIC960RxD Quad DTA, E1, 75 ohm unbalanced-ungrounded, 1.8 m 4-port cable
2877	ARTIC960RxD Quad DTA, H.100, 4-drop cable

9.5.3.3 Software Requirements

- AIX 4.2.1, or later
- DirectTalk for AIX, Version 2.1 LPP (5765-B81), to provide voice processing
- Adapter device driver (provided with adapter)

9.6 WAN Adapters

The following sections discuss the available WAN adapters.

9.6.1 2-Port Multiprotocol X.25 Adapter (# 2962)

The 2-Port Multiprotocol Adapter (# 2962) provides high-speed connections between stand-alone system units on a wide area network (WAN). This adapter adheres to the PCI standard and also supports SDLC and X.25 protocols. The 2-port Multiprotocol Adapter connects to WAN lines through externally attached data communication equipment including Channel Service Units (CSU), Data Service Units (DSU), and synchronous modems.

This adapter operates at speeds up to 2.048 Mbps and provides two ports that accommodate four selectable interfaces. These interfaces are:

- EIA 232D/V.24
- V.35
- V.36/EIA 449
- X.21

Interface configuration is selected by the type of cable attached. These cables are ordered separately, and you may configure the 2-Port Multiprotocol Adapter with two different cables.

9.6.1.1 Feature Characteristics

- 32-bit Bus Master PCI 2.1 Adapter
- Provides two, 36-pin, high-density (male) ports
- Provides four interface types: EIA 232D/V.24, V.35, V.36/EIA 449, and X.21
- Simultaneously supports two different interfaces
- Supports SDLC and X.25 full duplex synchronous protocols

9.6.1.2 Customer-Supplied Components

If you plan to operate this adapter using X.25 protocols, you must separately order the IBM AIXLINK/X.25 LPP (5696-926). This package provides a V.24, V.35, or X.21 port connection to X.25 packet-switched networks.

The system interface is determined by the cable connected to this adapter as listed in Table 108.

Table 108. Cable Information for the 2-Port Multiprotocol PCI Adapter

Cable Feature Code	Interface Configuration	Cable Terminations (Length)
2951	EIA 232D/V.24	36-pin to male DB25 (3 m)
2952	V.35	36-pin to 34-pin male (3 m)
2953	V.36/EIA 449	36-pin to 37-pin male (3 m)
2954	X.21	36-pin to male DB15 (3 m)

9.6.1.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, for SP nodes
- SDLC protocol support provided as part of the AIX Base Operating System
- X.25 protocol support requires a separately ordered LPP, IBM AIXLINK/X.25 (5696-926)

9.7 Asynchronous Transfer Mode Adapters

ATM is an international standard for high speed, cell relay, networking. ATM stems from its promise of high-speed networking, architected to transport any mixture of voice, video, and traditional computer data across the local, metropolitan, and wide area networks (LANs, MANs, and WANs).

ATM can use the network cabling currently installed. It can be used with voice or data grade Unshielded Twisted Pair (UTP) wiring (UTP categories 3 and 5, respectively) as well as cabling for other networking cable plans, such as Shielded Twisted Pair (STP) and fiber optic cable. This means that users will not have to face costly rewiring to adopt ATM networking technology.

ATM supports a wide range of transmission speed end environments and has no inherent speed limitations. The design of the interlace and lower layers of the protocol model is the same for LAN, MAN, and WAN networking environments with no predefined distance limits. Only the type of data being transmitted (based upon the ATM adaptation layer service) effects the class of networking services provided by the network.

The following is a list of the currently supported ATM adapters:

- Turboways 155 PCI UTP ATM Adapter (# 2963)
- Turboways 155 PCI MMF ATM Adapter (# 2988)
- Turboways 25 ATM Adapter (# 2998)

For additional information, refer to *RS/6000 ATM Cookbook*, SG24-5525.

9.7.1 Turboways 155 PCI UTP ATM Adapter (# 2963)

The Turboways 155 PCI UTP ATM Adapter (# 2963) enables TCP/IP applications to work in an asynchronous transfer mode (ATM) environment. This adapter provides dedicated 155 Mbps, full-duplex connection to ATM networks using either Permanent Virtual Circuits (PVC) or ATM Forum compliant Switched Virtual Circuits (SVC) UNI 3.1 signalling. TCP/IP supports PVC operation, but LAN Emulation does not. The adapter supports AAL-5 adaptation layer interface and communication with devices located on an ATM network, bridged token ring, Ethernet, or other LAN. LAN Emulation (LANE) is provided by the AIX operating system.

The Turboways 155 UTP ATM Adapter requires customer provided CAT5 High Speed Unshielded Twisted Pair (UTP) or Shielded Twisted Pair (STP) cables.

These cables must be certified for ATM operation. Maximum cable length is 100 m, and all cables must be terminated with RJ45 connectors.

9.7.1.1 Feature Characteristics

- 32-bit Bus Master PCI 2.1 adapter
- External RJ45 connector
- Provides signaling channel set up
- Provides virtual connection set up and tear down
- Supports point-to-point and point-to-multipoint switching
- Supports virtual circuits (maximum 1024)
- Supports classical IP and ATRP over ATM (RFC 1577)
- Supports Ethernet LAN Emulation and token ring
- Supports ATM SNMP

9.7.1.2 Customer-Supplied Components

You must supply the following components with this feature:

- Category 5 High Speed Unshielded Twisted Pair cables (or shielded) with RJ45 connectors (100 m maximum length).
- If you plan to use multipoint connections, you must provide an ATM switch.

9.7.1.3 Software Requirements

- AIX 4.3.2, or later
- PSSP 3.1, or later, on SP nodes

9.7.2 Turboways 155 PCI MMF ATM Adapter (# 2988)

The Turboways 155 PCI MMF ATM Adapter (# 2988) enables TCP/IP applications to work in an asynchronous transfer mode (ATM) environment. This adapter provides dedicated 155 Mbps, full-duplex connection to ATM networks using either Permanent Virtual Circuits (PVC) or ATM Forum-compliant Switched Virtual Circuits (SVC) UNI 3.1 signalling. The adapter supports AAL-5 adaptation layer interface and communication with devices located on an ATM network, bridged token ring, Ethernet, or other LAN. LAN Emulation (LANE) is provided by the AIX operating system.

9.7.2.1 Feature Characteristics and Requirements

- Provides signaling channel set up
- Provides virtual connection setup and tear down

- Supports point-to-point and point-to-multipoint switching
- Supports virtual circuits (maximum 1024)
- Supports classical IP and ATRP over ATM
- Supports Ethernet LAN Emulation and token ring
- Supports ATM SNMP

9.7.2.2 Customer-Supplied Components

- Plenum rated 62.5/125um multimode fiber cables terminated with an SC connector
- An ATM switch

9.7.2.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later

9.7.3 Turboways 25 ATM Adapter (# 2998)

The Turboways 25 ATM PCI Adapter (# 2998) is for direct access to ATM networks. This adapter provides dedicated 25 Mbps full-duplex connection using PVCs or SVCs and enables TCP/IP to run over an ATM network. The Turboways 25 ATM PCI Adapter also supports communication with devices located on an ATM network or bridged to a token ring, Ethernet, or other LAN.

9.7.3.1 Software Requirements

- AIX 4.1.5, 4.2.1, 4.3, or later

9.8 Ethernet and Token-Ring Adapters

The following section discusses the available Ethernet and token-ring adapters.

9.8.1 10/100 Ethernet 10BaseTx PCI Adapter (# 2968)

The IBM 10/100 Ethernet Tx PCI Adapter (# 2968) is a 10/100 PCI Ethernet adapter that is compatible with IEEE 802.3 and 802.3u specifications. The adapter has one RJ-45 connection that supports connections to 100BaseTx and 10BaseT networks.

9.8.1.1 Feature Characteristics

- Compatible with IEEE 802.3 standards
- 32-bit Bus Master PCI Bus 132 MBps

- Supports auto-negotiation of media speed and duplex operation
- Supports both full and half duplex operation over 10BaseT networks using the RJ-45 connector

9.8.1.2 Feature Components

- Adapter card
- Diskette with adapter device driver
- Installation instructions

9.8.1.3 Customer-Supplied Components

The customer must supply the following components for this feature:

- Network equipment, such as a Hub or Switch, required to attach to 10BaseT Ethernet LANs
- All Ethernet cables

Note

For 100BaseTX connections, Unshielded Twisted Pair (UTP) Category 5 cabling is required.

9.8.1.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, for SP nodes
- Adapter device driver (provided with adapter)

9.8.2 Gigabit Ethernet - SX Adapter (# 2969)

The PCI Gigabit Ethernet - SX Adapter (# 2969) is a 1000 Mbps PCI Ethernet adapter that is compatible with IEEE 802.3z specifications. The adapter has one external fiber connection that attaches to 1000BaseSX networks using 50 and 62.5 micron multimode cables with SC connectors.

This adapter will perform best in a 64-bit 50 MHz or 66 MHz slot, but will also function in a 32-bit 33 MHz slot.

9.8.2.1 Feature Characteristics

- Compatible with IEEE 802.3z Standards
- Supports full duplex operation over 1000BaseSX networks
- Supports jumbo frames with AIX 4.3.2 device driver, or later
- Works in 32- or 64-bit slots

9.8.2.2 Feature Components

- Adapter card
- Fiber wrap plug
- Installation instructions

9.8.2.3 Customer-Supplied Components

The customer must supply the following components for this feature:

- Network equipment, such as a switch or router, is required to attach to 1000BaseSX networks
- All Ethernet cables
- The maximum operating distances for the fiber cables are:
 - 260 meters with 62.5 micron multimode fiber
 - 440 meters with 50 micron multimode fiber

9.8.2.4 Software Requirements

- AIX 4.3.2, or later
- PSSP 3.1, or later, on SP nodes

9.8.3 PCI Ethernet BNC/RJ-45 Adapter (# 2985)

The PCI Ethernet BNC/RJ-45 Adapter (# 2985) is a 10 Mbps PCI Ethernet Adapter that is compatible with IEEE 802.3 specifications. The adapter has two external connections: BNC to attach to 10Base2 networks and RJ-45 to attach to 10BaseT networks.

9.8.3.1 Feature Characteristics

- 10 Mbps Ethernet compatible with IEEE 802.3 Standards
- 32-bit Bus Master PCI Bus 132 MBps
- Supports half duplex operations over 10Base2 networks using the BNC connector
- Supports both full and half duplex operation over 10BaseT networks using the RJ-45 connector

9.8.3.2 Feature Components

- Adapter card
- RJ-45 and AUI diagnostic wrap plugs
- Installation instructions

9.8.3.3 Customer-Supplied Components

- Network equipment, such as a Hub or Switch, required to attach to 10BaseT Ethernet LANs
- All Ethernet cables

Note

For 10BaseT connections, Unshielded Twisted Pair (UTP) Category 3, 4, or 5 cabling is required. UTP Category 5 cabling is strongly recommend to facilitate upgrades to 100 Mbps Ethernet LAN without cabling changes.

9.8.3.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver (part of base AIX BOS code)

9.8.4 PCI Ethernet AUI/RJ-45 Adapter (# 2987)

The PCI Ethernet AUI/RJ-45 Adapter (# 2987) is a 10 Mbps PCI Ethernet adapter that is compatible with IEEE 802.3 specifications. The adapter has two external connections: AUI to attach to 10Base5 networks and RJ-45 to attach to 10BaseT networks.

9.8.4.1 Feature Characteristics

- 10 Mbps Ethernet compatible with IEEE 802.3 Standards
- 32-bit Bus Master PCI Bus 132 MBps
- Supports half duplex operations over 10Base5 networks using the AUI connector
- Supports both full and half duplex operation over 10BaseT networks using the RJ-45 connector

9.8.4.2 Feature Components

- Adapter card
- RJ-45 and AUI diagnostic wrap plugs
- Installation instructions

9.8.4.3 Customer-Supplied Components

- Network equipment, such as a Hub or Switch, required to attach to 10BaseT Ethernet LANs
- All Ethernet cables

Note

For 10BaseT connections, Unshielded Twisted Pair (UTP) Category 3, 4, or 5 cabling is required. UTP Category 5 cabling is strongly recommend to facilitate upgrades to 100 Mbps Ethernet LAN without cabling changes.

9.8.4.4 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver (part of base AIX BOS code)

9.8.5 4-Port 10/100 Base-Tx Ethernet PCI Adapter (# 4951)

The IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter for RS/6000 (# 4951) makes available four Ethernet ports using a single PCI slot. This adapter provides the same function as purchasing four Ethernet adapters (# 2968).

The IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter for RS/6000 is a 32-bit long PCI adapter supporting four industry-standard Ethernet 10Base-T or 100Base-T interfaces supporting 10 or 100 Mbps data rates, either half or full duplex, on each of four separate ports. Each port is provided with an RJ-45 connector for attachment to standard CAT-3/5 Unshielded Twisted Pair, UTP cable. The adapter is IEEE 802.3u compatible and provides full auto-negotiation for detecting speed and duplex capability across each port.

The IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter for RS/6000 will provide network boot capability for selected RS/6000 machine types and models.

9.8.5.1 Feature Characteristics

- Auto negotiation of 10 Mbps or 100 Mbps
- Supports four RJ-45 connectors for UTP or STP cabling
- PCI Universal connector to accommodate +5 volt and +3.3 volt signaling
- Fits in full-size PCI slots
- Supports 32/64-bit PCI data width
- Operates at PCI bus speed of 33 MHz
- Supports half or full duplex operation
- Is FCC Class A and CISPR Class A certified for UTP-3/5 cabling
- Supports NIM Install on selected machines

9.8.5.2 Feature Components

- Adapter card
- Device driver for adapter card
- Firmware update
- Installations instructions

9.8.5.3 Customer-Supplied Components

The customer must supply the following components for this feature:

- Network equipment, such as a Hub or Switch, required to attach to 10/100 Mbps Ethernet LANs

- All Ethernet cables
- A standard CAT-3/5 Unshielded Twisted Pair, UTP cable

9.8.5.4 Software Requirements

- AIX Version 4.3.3 with the 2/2000 update CD, or later
- Adapter device driver
- Firmware update

Note

A firmware update can be acquired from the RS/6000 support page on the Internet or from your service team. The Internet address is:

<http://www.rs6000.ibm.com/support/micro/>

9.8.6 4/16 Mbps Token-Ring Adapter (# 4959)

The IBM token-ring PCI adapter for RS/6000 (# 4959) is a single slot, short, 32-bit PCI adapter supporting 4 and 16 Mbps data rates, either half-duplex or full-duplex. Automatic ring-speed selection is designed to prevent wrong speed insertion into the ring, even when connected to speed-sensing switches. This adapter will operate with either unshielded twisted pair (UTP) Cat. 5 cable with RJ-45 connectors or shielded twisted pair (STP) Type 1A cabling with 9-pin D-shell connectors. The token-ring PCI adapter will provide network boot capability.

9.8.6.1 Feature Characteristics

- Communicates at 4 or 16 Mbps
- Supports both UTP-5 (RJ45) and STP (DB9) cable connections
- Autodetects connection types at all speeds
- Meets PCI 2.1 Spec
- Fits in PCI half size slots
- Operates in 64-bit PCI slots in 32-bit mode
- Operates at PCI bus speed from 16 MHz to 33 MHz
- Supports full-duplex LAN operation at all speeds
- Consumes less than 2 watts of power
- Includes adapter and ring status LEDs
- Supports field update of FLASH EEPROM

- Offers on-card diagnostics in microcode
- Is FCC Class B and CISPR Class B certified for STP and UTP cabling

9.9 PCI Storage Adapters

This section contains adapter descriptions for the RS/6000 system PCI-type storage adapters.

9.9.1 SCSI-2 Fast/Wide RAID Adapter (# 2493)

The SCSI-2 Fast/Wide RAID Adapter (# 2493) implements RAID level 0, 1, and 5 support for SCSI-2 attached disks.

Configuration management, RAID algorithms, and error recovery are handled by an on-board 403 PowerPC.

Three internal 16-bit SCSI-2 SE 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.

9.9.1.1 Feature Characteristics

- Supports RAID 0, 1, and 5
- Controller conforms to ANSI document X3T9.2/86-109 revision 10k (SCSI-2) and the PCI Local Bus Specification Rev 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

9.9.2 PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)

The RS/6000 PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494) is a non-bootable, high-performance Ultra2 SCSI RAID Adapter providing RAID 0, 1, or 5 capability and can address up to 45 16-bit SCSI-2 physical disk drives on three independent SCSI buses. To increase the data writing performance, a 32 MB fast-write cache is provided as a resident part of this adapter that utilizes non-volatile RAM. During the unlikely event of an PCI 3-Channel Ultra2 SCSI RAID Adapter failure, a replacement PCI 3-Channel Ultra2 SCSI RAID Adapter can be installed and the fast-write cache can be removed from the failing adapter and installed in the new adapter insuring data integrity. The 32 MB fast-write cache can provide a significant improvement in data throughput and response time during certain sequence write operations compared to SCSI RAID adapters without the fast-write cache. The response time and data transfer improvement will vary depending upon the data block sizes, the percentage of sequential writes, machine type/model, and application parameters.

The PCI 3-Channel Ultra2 SCSI RAID Adapter has one internal and two external independent Ultra2 SCSI buses. The internal port can be used to provide an internal RAID solution on supporting RS/6000 systems. Systems with internal 6-pack disks can attach one PCI 3-Channel Ultra2 SCSI RAID adapter per internal 6-pack. The two external ports provide connectivity to an IBM 2104-DL1 Expandable Storage Plus Drawer or 2104-TL1 Expandable Storage Plus Tower at up to 80 MBps SCSI bus data rate. The two external ports can also be connected to the IBM 7131-105 external Fast/Wide SCSI disk enclosure using an optional VHDCI to P Converter Cable (# 2118).

Note

Even though the PCI 3-Channel Ultra2 SCSI RAID Adapter has ports that run at Ultra2 SCSI speeds (up to 80 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system backplane. Also, the externally attached 7131-105 will only run at a maximum of Fast/Wide SCSI speeds (up to 20 MBps).

To take full advantage of the Ultra2 (up to 80 MBps) speed of this adapter, the proper AIX level also needs to be considered. AIX 4.2.1 with appropriate APAR updates and 4.3.3, or later, supports the full range of SCSI bus data rates (including Ultra2 SCSI up to 80 MBps).

AIX 4.3.2 only supports data rates up to 40 MBps. The appropriate AIX APAR needed are as follows:

- AIX 4.2.1 with APAR IX86638, IX86654, IX87183, IY01945, and IY02373
- AIX 4.3.2 with APAR IX86395 and IX86396

Note

It is also possible to upgrade an old PCI 3-Channel Ultra SCSI RAID Adapter (# 2494) to the newer PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494). To upgrade to the Ultra2 version (up to 80 MBps capability), a new level of microcode needs to be downloaded. To obtain the latest level of microcode, go to the following Internet address and follow the microcode download procedure:

<http://www.rs6000.ibm.com/support/micro/download.html>

For best overall performance and data protection, you should evenly distribute hard disk drives across the three available channels.

9.9.3 PCI Universal Differential Ultra SCSI Adapter (# 6204)

The IBM PCI Universal Differential Ultra SCSI Adapter is the latest technology advancement of an RS/6000 SCSI-2 differential adapter with a maximum data transfer rates of 40 MBps. This adapter has the capability to be plugged into the newer +3.3 volt PCI slots and the older +5 volt PCI slots. The PCI Universal Differential Ultra SCSI (# 6204) allows connection to external differential SCSI-2 F/W or Ultra SCSI type devices up to 25 meters away.

The adapter will negotiate with each external device and transfers data at the fastest SCSI data transfer rate capability of the external device.

The adapter conforms to SCSI-2 standard and the Fast-20 (Ultra) documentation. Industry-standard, 68-pin connectors are incorporated (SCSI "P" connector definition of X3T9.2/90-048).

9.9.3.1 Feature Characteristics

- Data transfer rates of 40 MBps
- Capable of plugging into +3.3 volt or +5 volt PCI slots
- Connection to external differential SCSI-2 F/W or Ultra SCSI
- Negotiates with each external device and transfers data at the fastest SCSI data transfer rate capable by the external device

- Industry-standard, 68-pin connectors

9.9.3.2 Software Requirements

- AIX Version 4.3.3, or later
- PSSP 3.1.1, or later, on the SP nodes

9.9.4 PCI Dual Channel Ultra2 SCSI Adapter (# 6205)

The PCI Dual Channel Ultra2 SCSI Adapter (# 6205) is a 64-bit adapter and is an excellent solution for high performance SCSI applications. The PCI Dual Channel Ultra2 SCSI Adapter provides two SCSI channels (busses). Each SCSI bus can either be internal or external and will support a data rate of up to 80 MBps, up to twice the maximum data transfer rate of previous Ultra SCSI adapters (40 MBps).

To achieve an Ultra2 SCSI bus data rate of up to 80 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage Differential (LVD) drivers and receivers. To utilize this Ultra2 80 MBps performance, all attaching devices or subsystems must also be Ultra2 LVD devices. If any device is not Ultra2 LVD, the adapter will switch its SCSI bus to single-ended (SE) performance and interface at the lower SE SCSI bus data rate of the device.

Two industry standard VHDCI 68-pin connectors are mounted on the adapter's end bracket allowing attachment of various LVD and SE external subsystems. A three meter converter cable, VHDCI to P, Mini-68 pin to 68 pin, (# 2118) can be used with older external SE subsystems to allow connection to the VHDCI connector on the PCI Dual Channel Ultra2 SCSI Adapter.

Any supported RS/6000 system can be set up to boot from the PCI Dual Channel Ultra2 SCSI Adapter (# 6205). If you are running with AIX 4.3.3, or later software, this adapter has native boot support as part of that level of AIX software. If you are running AIX 4.2.1 software, the following procedure applies to booting using the PCI Dual Channel Ultra2 SCSI Adapter:

- The designated boot SCSI disk can be located under the covers of a processor unit or in an external SCSI storage unit.
- AIX Version 4.2.1 must be loaded to the designated SCSI boot disk using AIX Network Install Manager (NIM) before booting from the SCSI boot disk.
- The system with a designated SCSI boot disk must have a network connection with another RS/6000 system performing the NIM Master

function to perform the install. On RS/6000 SP systems, a similar network install is performed from a control workstation.

- Once AIX, with updates, is installed on the designated SCSI boot disk and the system is configured for booting, booting will take place from the boot disk drive without any support from the control processor or NIM Master and the system does not have to be connected to the network at boot time.

9.9.4.1 Software Requirements

- AIX Version 4.2.1, 4.3.3, or later

Note

To connect a #6205 to the following devices:

- 7204-409 External 9.1 GB disk drive
- 7204-419 External 18.2 GB disk drive
- 7206-220 External 4 mm tape drive
- 7332-220 External 4 mm tape drive

You need to order the no-charge VHDCI Cable/Interposer (# 9799) and one of the appropriate cables:

- 9137 1.0 m Ultra2 SCSI cable
- 9149 1.0 m SCSI-2 F/W cable

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

The PCI Single-Ended Ultra SCSI Adapter (# 6206) provides a single-ended SCSI-2 Ultra/Wide interface that can burst data between devices on the SCSI bus at 40 MBps (twice the fast/wide rate) using block sizes greater than 64 KB. It conforms to SCSI-2 standards and Fast-20 (Ultra) documentation. The PCI Single-Ended Ultra SCSI Adapter supports both internal and external devices connected to the same SCSI bus. Industry standard SCSI P (68-pin) connectors are incorporated on the adapter.

9.9.5.1 Feature Characteristics

- 32-bit Bus Master PCI 2.1 Adapter
- Supports attachment of internal and external single-ended 8-bit and 16-bit SCSI or Ultra SCSI devices:
 - External connections on J2 with 68-pin SCSI-3 standard P connector

- Internal connections on J3 with 68-pin high-density SCSI connector for 16-bit attachments
- Internal connections on J4 with 50-pin (2x25) SCSI connector for 8-bit attachments

9.9.5.2 Adapter Limitations

- Data transfer rates are limited to the speed of the slowest attached device. For example, if you connect an Ultra drive and a fast/wide drive, the adapter will limit data transfers to fast/wide rates.
- If a cable is attached to the external J2 connector, data transfer rates will be limited to fast/wide rates.
- Ultra data transfer rates can only be achieved using the internal connections with cable lengths of 1.5 m or less.
- External cable lengths are limited to 3 m for fast/wide data transfer rates.
- The internal J3 and J4 connectors cannot be used at the same time.
- Single-ended (SE) SCSI Adapters cannot interoperate with Differential SCSI Adapters in twin-tailed (high availability) configurations.

9.9.5.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on the SP nodes

9.9.5.4 Cable Options

The following optional cables are available for the (# 6206) SCSI adapter:

- (# 2117) 16-bit SE external Y-cable, 0.9 m
- (# 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m
- (# 2425) 16-bit adapter-to-adapter SCSI cable, 2.5 m

All cables must conform to X3T9.2/90-048 standards.

Note

Due to the short length of PCI SCSI cables, you must pay close attention to cable planning.

9.9.6 PCI Differential Ultra SCSI Adapter (# 6207)

The PCI Differential Ultra SCSI Adapter (# 6207) provides a differential SCSI-2 Ultra/Wide interface that can burst data between devices on the SCSI

bus at 40 MBps. The PCI Differential Ultra SCSI Adapter supports Ultra and fast/wide synchronous data transfers, and it supports external devices (no internal connections) up to 25 m away. This adapter conforms to SCSI-2 standards and the Fast-20 (Ultra) documentation. Industry standard SCSI P (68-pin) connectors are incorporated on the adapter.

9.9.6.1 Feature Characteristics

- 32-bit Bus Master Adapter
- Supports attachment of external 8-bit or 16-bit SCSI devices on the J2 port using a 68-pin SCSI-3 standard connector

9.9.6.2 Adapter Limitations

- Data transfer rates with # 6207 are limited to the speed of the slowest device on the SCSI bus.
- Single-ended (SE) and double-ended SCSI adapters cannot be twin-tailed to the same external disk array when used in a high-availability configuration.

9.9.6.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes

9.9.6.4 Cable Options

The following optional cables are available for the PCI SCSI-2 Ultra/Wide Differential Adapter:

- (# 2114) 16-bit DE external Y-cable, 0.9 m
- (# 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m long
- (# 2425) 16-bit adapter-to-adapter SCSI cable, 2.5 m long

All cables conform to X3T9.2/90-048 standards.

9.9.7 SCSI-2 Single-ended Fast/Wide Adapter (# 6208)

The PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208) provides a single-ended SCSI-2 Fast/Wide interface that can burst data between devices on the SCSI bus at 20 MBps. It conforms to SCSI-2 standards and supports fast/wide synchronous data rates of up to 10 MHz. Feature # 6208 supports both internal and external devices connected to the same SCSI bus.

9.9.7.1 Feature Characteristics

- 32-bit Bus Master Adapter

- Supports attachment of internal and external single-ended 8-bit and 16-bit SCSI devices:
 - External connections on J2 with 68-pin SCSI-3 standard P connector
 - Internal connections on J3 with 68-pin high-density SCSI connector for 16-bit attachments and on J4 with 50-pin SCSI connector for 8-bit attachments

9.9.7.2 Feature Limitations

- The J3 and J4 connectors cannot be used at the same time.
- Supports a maximum cable length of 25 m.
- Single-ended (SE) SCSI adapters cannot interoperate with Differential SCSI adapters in twin-tailed (high availability) configurations.

9.9.7.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes

9.9.7.4 Cable Options

The following optional cables are available for the PCI SCSI-2 Fast/Wide Single Ended Adapter:

- (# 2117) 16-bit SE external Y-cable, 0.9 m
- (# 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m
- (# 2425) 16-bit adapter-to-adapter SCSI cable, 2.5 m

All cables conform to X3T9.2/90-048 standards.

Note

Due to the short length of PCI SCSI cables, you must pay close attention to cable planning.

9.9.8 SCSI-2 Differential Fast/Wide Adapter (# 6209)

The PCI SCSI-2 Fast/Wide Differential Adapter (# 6209) provides a differential SCSI-2 Fast/Wide interface that can burst data between devices on the SCSI bus at 20 MBps. It conforms to SCSI-2 standards and supports fast/wide synchronous data rates of up to 10 MHz. SCSI-2 Differential Fast/Wide Adapter 4-B supports external devices connected to the same SCSI bus.

9.9.8.1 Feature Characteristics

- 32-bit Bus Master Adapter
- Supports attachment of external 16-bit SCSI devices on the J2 port using a 68-pin SCSI-3 standard P connector

9.9.8.2 Feature Limitations

- Single-ended (SE) and double-ended SCSI Adapters cannot be twin-tailed to the same external disk array when used in a high-availability configuration.

9.9.8.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes

9.9.8.4 Cable Options

- (# 2114) 16-bit DE external Y-cable, 0.9 m
- (# 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m
- (# 2425) 16-bit adapter-to-adapter SCSI cable, 2.5 m

All cables must conform to X3T9.2/90-048 standards.

Note

Due to the short length of PCI SCSI cables, you must pay close attention to cable planning.

9.9.9 SSA RAID 5 Adapter (# 6215), SSA Fast-Write (# 6222)

The PCI SSA RAID 5 Adapter (# 6215) supports RAID 5 SSA disk arrays and can be used to access non-RAID disks (JBOD) between multiple hosts. It has the capability to improve write response time in the single initiator mode for both RAID and non-RAID disks by the addition of the Fast-Write Cache Option (# 6222).

The SSA Fast-Write Cache is an optional 4 MB fast-write module that plugs into the PCI SSA RAID 5 Adapter (# 6215). The cache option (# 6222) uses nonvolatile RAM having over seven years of memory retention. Nonvolatile memory allows you to transfer the cache module from a failing multi-initiator adapter to a new adapter during the unlikely event of an adapter failure. This helps insure data integrity and operational reliability. The fast-write module is not supported for multi-adapter loops.

9.9.9.1 Devices Supported

- All 7133 models
- 7131-405 internal RS/6000 SSA disk drive configurations
- The NAPA Optical Extender

9.9.9.2 Feature Characteristics and Requirements

- 32-bit PCI bus.
- Support for floating hot spares on the same loop.
- RAID five arrays from (2+P) up to (15+P).
- Up to six (15+P) or 32 (2+P) RAID 5 array groups per adapter.
- Up to eight adapters are supported in one loop for JBOD.
- Up to two adapters are supported in one loop for RAID 5.
- Only one Fast-Write Cache module is supported on each PCI SSA RAID 5 Adapter.
- All members of a RAID 5 array must be on the same SSA loop.
- When more than one SSA RAID Adapter is used, the adapters should be distributed evenly across all available PCI slots.
- The MCA SSA RAID Adapter (# 6219) may be used on the same loop.

9.9.9.3 Software Requirements

- AIX 4.2.1, or later
- PSSP 2.4, or later, on SP nodes
- Adapter device driver and FDDI common code (provided with adapter)

9.9.10 Advanced SerialRAID Plus Adapter (# 6230)

The Advanced SerialRAID Plus Adapter (# 6230) is a 4-port (2 loop) Serial Storage Architecture (SSA) adapter providing an instantaneous data transfer rate of up to 160 MB per loop.

It provides the following features:

- 1- or 2-way RAID 0+1 with 4-16 disks (Up to eight mirrored pairs of disks per array)
- 1- or 2-way RAID 1 with two disks per array
- Optional 32 MB Fast Write Cache (# 6235) for dual or single attach configurations

- Optional 128 MB DRAM (# 6231) (vs. standard 64MB DRAM) for full use of 32 MB cache in 2-way operation

In addition, these functions were already supported on the 6225 and are now also supported on the 6230:

- Up to 8-way non-RAID support.
- Up to 2-way RAID 5.
- 1-way RAID 0.
- Remote site mirroring up to 10 km with advanced optical extenders.
- SSA Boot support.
- Hot spares are supported with RAID 0+1 arrays as well as RAID 5 arrays.

The IBM 7133 Serial Disk System and the Advanced SerialRAID Plus Adapter provide an exceptional storage solution for RS/6000 servers. The new RAID 1 enhancement provides data protection via hardware mirroring, while the RAID 0+1 enhancement provides hardware mirroring data protection and the added performance of data striping.

With the new RAID 0+1 function, the Advanced SerialRAID Plus adapter can provide significantly better write response time for database applications than was possible with AIX Logical Volume mirroring. With AIX LVM mirroring, mirror write consistency was generally required to be active when used with database applications. Mirror write consistency isn't usually needed when the RAID 1 function is provided by the Advanced SerialRAID Plus Adapter hardware – allowing, in some cases, write response time to be cut in half when compared to that of AIX LVM mirroring with mirror write consistency.

Also, the RAID 0+1 function provides a performance boost over RAID 5 for writes and 70/30 random workload performance. The RAID 0+1 performance measurements with the Advanced SerialRAID Plus Adapter include:

- 1-way: 5000 I/Os per sec 70/30 Random 4 Kb I/Os/sec, 50 MBps Write data rate, 80 MBps Read data rate.
- 2-way: 7500 I/Os per sec 70/30 Random 4 Kb I/Os/sec, 65 MBps Write data rate, 150 MBps Read data rate.

Note

These are lab measurements and may not be realized in all customer environments.

Since December 17, 1999, the new microcode and device driver code, which enables the 2-way Fast Write Cache, RAID 1, and RAID 0+1 function for existing Advanced SerialRAID Adapters (# 6225), is available free of charge. This function for the Advanced SerialRAID Adapter (# 6225) can be received by downloading the code, documentation, and installation instructions from the following Web site:

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

Note

This new microcode and device driver code is not being shipped from IBM with the Advanced SerialRAID Adapter (# 6225). It can only be downloaded from the above Web site for that adapter.

A Q&A document on the November 23, 1999, SSA announcements is attached here.

The Advanced SerialRAID Plus Adapter # 6230 supports 7133 models 010, 020, 500, 600, D40, and T40, in addition to providing support for under-the-cover SSA RS/6000 implementations.

Note

The Advanced SerialRAID Plus Adapter # 6230 can coexist on the same loop with the Advanced SerialRAID adapter # 6225, as well as with the SSA RAID EL adapters # 6215 and # 6219. Note that 2-way Fast-Write Cache and many other capabilities are not supported with # 6215 and # 6219.

Additional information, a list of the supported servers for the IBM 7133 and the Advanced SerialRAID Plus Adapter, as well as an updated Spec sheet for the RS/6000 SSA Adapters can be found at the following Web site:

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

The adapter accepts a 32 MB Fast-Write Cache Option Card (# 6235) that improves write performance in RAID 0+1, RAID 5, and non-RAID applications. When the 32 MB Fast-Write Cache Option Card is used, the adapter can be configured in either single or dual initiator fast-write cache mode. In dual initiator fast-write cache mode, if one of the two adapters fails, the failing adapter is designed to transfer control over to the other. Also, in dual initiator fast-write cache mode, the 128 MB DRAM Option Card (# 6231) is available and required to utilize the full 32 MB of fast-write cache on the adapters. If the 128 MB DRAM Option Card is not used in dual initiator

fast-write cache mode, the effective fast-write cache capacity will be 16 MB on each adapter.

When operated in a RAID 5 configuration, the Advanced SerialRAID Plus Adapter will support (2+P) to (15+P) arrays and up to six (15+P) arrays. When operated in a RAID 1 or RAID 0+1 configuration, it will support up to eight mirrored disk drives. The adapter also supports hot spares in RAID 5 and RAID 0+1 mode. The Advanced SerialRAID Plus Adapter also supports connectivity to external disk enclosures and internal RS/6000 SSA configurations. An optional SSA Fiber-Optic Extender is also supported.

Any supported RS/6000 system can be set up to boot from an Advanced SerialRAID Plus Adapter (# 6230), provided a non-RAID SSA disk is included as part of the configuration. Other disks associated with the adapter can be RAID but at least one disk must be a non-RAID SSA disk. The non-RAID SSA disk can be located under the covers of a processor unit or in an external SSA storage unit. If your system is running AIX 4.3.3, or later, software, native boot capability is supported. For factory system orders with AIX preload requested, an internal SCSI disk drive will be preloaded as the native boot disk even if internal SSA disk drive(s) are present. If your system is running AIX 4.2.1 or AIX 4.3.2 software, the following procedure applies to booting using the Advanced SerialRAID Plus Adapter:

- The non-RAID SSA disk can be located under the covers of a processor unit or in an external SSA storage unit.
- A supported AIX version of software (with proper updates) must be loaded to the non-RAID SSA disk using the AIX Network Install Manager (NIM) before booting from the non-RAID disk.
- The system with a non-RAID SSA disk must have a network connection with another RS/6000 system performing the NIM Master function to perform the install. On RS/6000 SP systems, a similar network install is performed from a control workstation.
- Once AIX (with updates) is installed on the non-RAID SSA disk and the system is configured for booting, booting will take place from the boot disk without any support from the control processor or NIM Master and the system does not have to be connected to the network at boot time.

When performing SSA device or subsystem planning, installation, upgrades, or preventive maintenance, refer to the following Web support pages. They provide access to the latest SSA publications and support code for the system, SSA adapters, and SSA subsystems.

The following sites should also be considered during system hardware or operating system upgrades if SSA devices are included in the system.

<http://www.hursley.ibm.com/ssa/>

Contains links to SSA publications and other SSA Web sites, including the following:

<http://www.hursley.ibm.com/ssa/rs6k/>

Note

Internal ports on the adapter are not supported. See machine/model-specific information to determine if internal SSA disk drives and associated hardware/cables are supported.

There has been some confusion in the past about the use of SSA Fibre Optical Extenders. Table 109 shows an overview of the compatibility and use of both Optical Extender generations.

Table 109. Overview of the Compatibility and Use of Both Optical Extenders

Optical Extender Type		SSA Fibre Optical Extender	SSA Advance Fibre Optical Extender
Feature Code		# 5500	# 8851
Link Speed	20 MBps	yes	yes
	40 MBps	not supported	yes (only with Advanced SerialRAID Adapter # 6225 and 7133 Model D40/T40)
Fibre Cable Type	multimode	either 50 μ m or 62.5 μ m	either 50 μ m or 62.5 μ m (Mode Conditioning Patch Cords # 8852 or # 8853 must be installed)
	singlemode	not supported	9 μ m
Max. Distance	multimode	up to 2.4 km	up to 3 km with # 6225 or # 6230, up to 2.4 km with all other adapter types

Optical Extender Type		SSA Fibre Optical Extender	SSA Advance Fibre Optical Extender
	singlemode	not supported	up to 10 km (only with Advanced SerialRAID adapter # 6225 or # 6230, and 7133 Model D40/T40)
Max. Overall Attenuation		3 dB	8 dB
Supported drawers		all models	all models
Supported Adapter Types		all	all

9.9.11 Gigabit Fibre Channel Adapter (# 6227)

The Gigabit Fibre Channel Adapter Adapter (# 6227) provides single initiator capability over an optical fiber link (loop) running up to 100 MBps.

Two initiator capability is available when the RS/6000 Gigabit Fibre Channel Adapter is attached to the IBM Fibre Channel RAID Storage Server (2102-F10), also at 100 MBps. With the use of optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage.

Note

Gigabit Fibre Channel Adapter for PCI Bus (#6227) only supports attachment to the IBM 2102-F10 Fibre Channel RAID Storage Server with IBM 2102-D00 Expandable Storage Units through optional IBM 2103-H07 Fibre Channel Storage Hubs. Attachment support through the IBM 2108 Storage Area Network (SAN) Data Gateway Model G07 to specific storage subsystems is also supported.

Refer to the following IBM storage subsystem Web page for additional supported server attachment information:

<http://www.storage.ibm.com/hardsoft/disk/products.htm>

9.9.11.1 Software Requirements

- AIX Version 4.3.2, or later

9.10 Hot-Pluggable Adapters

In February 2000, IBM announced new capability for PCI adapters: Hot Plug PCI Adapters. This function enable maintenance of these adapters without powering off and rebooting the system. Hot Plug PCI Adapters capability is currently supported with 7025-F80, 7026-H80, 7026-M80, and SP High nodes (only in the SP I/O Expansion Unit). Table 110 lists the hot-pluggable PCI Adapters.

If you are running HACMP on a system that allows Hot-Plug removal and replacement of adapters, special care must be taken when replacing network adapters that are part of an active HACMP cluster configuration.

Table 110. Hot Pluggable PCI Adapters

Feature Code	PCI Adapter Name	F80	H80	M80	SP I/O Expansion Unit
2493	SCSI-2 F/W RAID	info N/A	info N/A	info N/A	info N/A
2494	PCI 3-Channel Ultra2 SCSI RAID Adapter	Y	Y	Y	info N/A
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	Y	Y	Y	info N/A
2732	IBM Short-wave Serial HIPPI PCI Adapter for RS/6000	info N/A	info N/A	info N/A	Y
2733	IBM Long-wave Serial HIPPI PCI Adapter for RS/6000	info N/A	info N/A	info N/A	Y
2741	SysKonnnect SK-NET FDDI-LP SAS PCI for PowerPC Systems	Y	Y	Y	Y
2742	FDDI SK-NET LP DAS	Y	Y	Y	Y
2743	FDDI SK-NET UP SAS	Y	Y	Y	Y
2751	S/390 ESCON Channel Adapter	N	N	N	Y
2943	8-Port Asynchronous Adapter EIA 232/RS-422	Y	Y	Y	Y
2944	128-Port Asynchronous Adapter EIA-232	Y	Y	Y	Y
2947	ARTIC960Hx 4-Port Selectable Adapter	N	N	N	N
2948	ARTIC960Hx 4-Port T1/E1 PCI	info N/A	info N/A	info N/A	info N/A
2962	2-port Multiprotocol X.25 Adapter	Y	Y	Y	Y
2963	ATM 155 Turboways UTP Adapter	Y	Y	Y	Y
2968	Ethernet 10/100 MB	Y	Y	Y	Y
2969	Gigabit Ethernet - SX	Y	Y	Y	Y

Feature Code	PCI Adapter Name	F80	H80	M80	SP I/O Expansion Unit
2985	PCI Ethernet BNC/RJ-45 Adapter	Y	info N/A	info N/A	Y
2987	PCI Ethernet AUI/RJ-45 Adapter	info N/A	info N/A	info N/A	Y
2988	ATM 155 Turboways MMF Adapter	Y	Y	Y	Y
2998	ATM 25 Turboways	info N/A	info N/A	info N/A	info N/A
4951	4-Port 10/100 Base-TX Ethernet	info N/A	info N/A	info N/A	info N/A
4959	4/16 Mbps token-ring adapter	Y	Y	Y	Y
6204	PCI Universal Differential Ultra SCSI	Y	Y ²	Y ²	Y
6205	PCI Dual Channel Ultra2 SCSI Adapter	Y	Y ²	Y ²	Y
6206	Ultra SCSI Single-Ended	info N/A	info N/A	info N/A	Y
6207	Ultra SCSI Differential	Y	info N/A	Y	Y
6208	SCSI-2 F/W Single-Ended	Y	Y ²	info N/A	info N/A
6209	SCSI-2 F/W Differential	info N/A	info N/A	info N/A	info N/A
6215	SSA RAID 5 (accepts optional SSA Fast-Write Cache module (# 6222))	info N/A	N	info N/A	info N/A
6227	Gigabit Fibre Channel Adapter	Y	Y	Y	Y
6230	Advanced SerialRAID Plus Adapter	Y ²	Y ²	Y ²	Y
6310	ARTIC960RxD Quad Digital Trunk Adapter	Y ¹	Y ¹	Y	info N/A
6311	IBM ARTIC960RxF PCI Adapter for RS/6000	N	Y	Y	info N/A
8396	IBM RS/6000 SP System Attachment Adapter	info N/A	info N/A	info N/A	info N/A
¹ This adapter is hot-pluggable if it is not cabled to another card internally. ² Do not hot-plug any PCI adapter supporting the system's boot device or system console.					

Chapter 10. Graphics Accelerators

IBM offers a broad range of versatile graphics accelerators to meet your 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 an industry-leading 2D accelerator 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 IBMs family of 2D graphics accelerators also 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, such as OpenGL. These graphics accelerators are designed to off load the system processor by providing varying amounts of hardware acceleration.

From a 3D graphics perspective, IBMs 3D graphics accelerators fall into three basic hardware classes:

- Class I** 2D and entry-level 3D
All computing and rasterization work is performed on the CPU. 3D capabilities are achieved through use of the SoftGraphics product.
- Class II** Mid-range 3D
These accelerators provide hardware acceleration support for anti-aliasing, texture mapping, and rasterization. Geometry processing (lighting and transforms) is done on the CPU.
- Class III** High-end 3D
Hardware acceleration for rasterization and geometry processing is offloaded to DSPs on the accelerator that, in turn, makes use of custom rasterization chips.

In this chapter, the graphics accelerators are discussed and a performance table of each graphics accelerator is presented. When examining the performance tables, there are a few definitions that need to be explained.

Values shown here are the results of development-level systems. All performance data was obtained in a specific environment and is presented as

is for illustrative purposes only. While these values should be indicative of generally-available systems, no warranties or guarantees are stated or implied by IBM. IBM recommends application-oriented testing for performance predictions and offers these commonly reported benchmarks only as an initial indicator. Additional information on these tests is available from your IBM local Branch Office or IBM Authorized Reseller. Additional information can also be obtained on the World Wide Web at:

<http://www.specbench.org>

The following GPC benchmarks reflect the performance of the microprocessor, memory subsystem, and graphics accelerator:

GPC/XPC results

Xmark93 is the weighted geometric mean of 447 tests executed in the x11perf suite and is an indicator of 2D graphics performance in an X environment. Larger values indicate better performance.

GPC/PLB results (graphics)

PLBwire93 and PLBsurf93 are geometric means of literal and optimized Picture Level Benchmark (PLB) tests for 3D wireframe and 3D surface tests, respectively. The benchmark and tests were developed by the Graphics Performance Characterization (GPC) Committee. The results shown used the graPHIGS API. Larger values indicate better performance.

GPC/OPC results

DX-05, DRV-06, Light-03, AWadv-02, and Pro-CDRS-02 are weighted geometric means of individual viewset metrics. The viewsets were developed by Independent Software Vendors (ISVs) with the assistance of OpenGL Performance Characterization (OPC) member companies. Larger values indicate better performance.

10.1 Available Graphics Accelerators

This section discusses in detail the available graphics accelerators and graphics accelerators that were lately withdrawn from the market.

The POWER GXT3000P and POWER GXT2000P are the latest graphics accelerator offerings in mid-range 3D and high-end 3D, providing significantly improved performance over the GXT800P and GXT500P accelerators they replaced. In 2D and entry-level 3D the POWER GXT300P and GXT130P replaced the GXT250P/GXT255P and GXT120P graphics accelerators.

The functional design of the GXT3000P provides local hardware support for lighting. This capability provides enhanced performance for those applications that use directional (infinite) light sources, but it requires more data to be passed across the bus to provide this function. Therefore, the performance of the GXT3000P is highly dependent on the system throughput of the host workstation. The combination of the processor, memory subsystem, and PCI bus structure found on the Model 150 does not allow the GXT3000P to perform in the optimal range. As a result, for most applications running on the Model 150, a GXT2000P will provide a level of performance similar of the GXT3000P. Therefore, the GXT2000P is the preferred graphics accelerator on this platform.

On the Model 170, 260, and 270, the system throughput is substantially greater than that of the Model 140 and 150, due to the POWER3 or POWER3-II processor design, the I/O backplane, and the 64-bit wide 50 MHz PCI bus slots. As a result, the GXT3000P provides enhanced performance for many applications, depending on the functions used by the application. If, as a part of everyday use, there is significant use of texture mapping or directional (infinite) light sources, the GXT3000P will provide greater performance than the GXT2000P. For users running applications, such as I-DEAS, CATIA, Pro/Engineer, and in-house MCAD software in the 3D solid, shaded, or lit rendering modes (for example, designers looking at solid surfaces), the GXT3000P may show up to twice the performance over the GXT2000P (when not fill bound). This is due to the GXT3000P support of hardware lighting and a higher polygon setup rate. In instances where users have large, complex models and geometry processing as a bottleneck, the GXT3000P will have a distinct advantage, again due to the higher polygon setup rate. Applications using texture mapping (for example, digital content creation, games, such as Quake, industrial design, and large model visualization) should see a two to three times advantage in performance over the GXT2000P. Therefore, the performance advantages of the GXT3000P over the GXT2000P are dependent upon the application usage, size of model, and functional characteristics. The key to proper selection of the graphics accelerator is knowledge of the intended use of the workstation and the type of data and models to be used.

10.1.1 POWER GXT3000P Graphics Accelerator (# 2825)

The POWER GXT3000P Graphics Accelerator for 3D visualization marks a breakthrough in performance and function for design and visualization solutions. It is the highest performing IBM graphics accelerator available for RS/6000 workstations. The graphics subsystem, combined with IBMs POWER3-II, POWER3, and PowerPC-based workstations, delivers outstanding speed and 3D performance for demanding applications in

mechanical design and engineering, for aerospace and automotive, petroleum exploration and production, scientific visualization, and other technical industries.

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

The GXT3000P was designed to support the features of PHIGS and OpenGL 1.2 (including hardware texturing), plus many OpenGL extensions. It supports display resolutions of up to 1280 x 1024 and 1024 x 768 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 ISO 9241 Part 3-compliant components including monitors.

The GXT3000P graphics accelerator is a 64-bit adapter that requires a single card slot. It will fit in either a 64-bit or 32-bit slot.

On the Model 260 or 270, a 64-bit slot 2 is recommended. If installed in the 32-bit slot 4 on the Model 260 or 270, slot 3 must remain empty. A maximum of two accelerators are allowed.

On the Model 170, you can install it in slot 2 or 4 and a maximum of two accelerators can be installed.

On the Model 150, it will attach using the 32-bit slot; however, an adjacent slot must be left empty. On the Model 150, the GXT3000P must be installed in slot 3 and only one accelerator in each system is allowed.

Advanced features include support of:

- Up to eight infinite light sources in hardware
- 3D textures in hardware
- Stereo viewing
- Eight window ID bits or color tables
- Up to 10 million polygon vertices per second

The GXT3000P graphics accelerator is supported on AIX 4.2.1, AIX 4.3.2, or later.

10.1.1.1 Supported Machines

The following machines support the GXT3000P:

- 7043-150
- 7044-170
- 7043-260
- 7043-270 (Upgrade from 7043-260)
- 7044-270

Table 111 lists the supported and most applicable display types and cables for the POWER GXT3000P.

Table 111. GXT3000P Display and Cable Table

Display Type	Cable Feature Code
5081	N/S
6091 - 16	N/S
6091 - 19	N/S
6091 - 19i	# 4217
6091 - 23	N/S
POWERdisplay 16	N/S
POWERdisplay 17	# 4217
POWERdisplay 19	# 4217
POWERdisplay 20	# 4217
P50	Cable included with display
P70	# 4238
P72	Cable included with display
P76	Cable included with display
P92	Cable included with display
G52	N/S
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
T55A White	N/S

Display Type	Cable Feature Code
T55A Black	N/S
9516-B03	# 4217
9516-B04	# 4217
T85A White	Cable included with display
T85A Black	Cable included with display
Note: *N/S = This accelerator/display combination is not supported	

10.1.1.2 Performance Table

Table 112 and Table 113 provide the performance results of the POWER GXT3000P graphics accelerator.

Table 112. GXT3000P Performance Table - Historic Measures

Machine Description	Xmark 93	PLBwire 93	PLBsurf 93	CDRS- 04	DX- 04	DRV- 05	Light -02	Awadv -02	ProCDRS- 01
43P-150 (375 MHz) GXT3000P AIX 432	37.17	257.7	467.00	96.72	11.03	6.33	1.30	14.50	
43P-260 (200 MHz) GXT3000P AIX 432	37.35	436.4	598.10	248.05	15.90	7.51	1.21	12.87	
43P-260 (2X200 MHz) GXT3000P AIX 432	37.09	625.3	843.00	248.05	15.90	7.51	1.21	12.87	
43P-150 (375 MHz) GXT3000P AIX 432 - 3/99 Update CD	37.69	255.6	492.30	94.70	10.92	6.28	1.22	14.27	11.41
43P-260 (200 MHz) GXT3000P AIX 432 - 3/99 Update CD	37.18	435.1	652.50	245.46	16.66	7.36	1.13	12.82	24.97
43P-260 (2X200 MHz) GXT3000P AIX 432 - 3/99 Update CD	37.18	702.2	897.90	245.46	16.66	7.36	1.13	12.82	24.97

Table 113. GXT3000P Performance Table - Current Measures

Machine Description	Xmark 93	PLBwire 93	PLBsurf 93	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
43P-150 (375 MHz) GXT3000P AIX 4.3.3.1	36.2	254	498	10.5	6.14	1.14	15.7	11.4
44P-170 (333 MHz) AIX 4.3.3.1 L2 - 1MB	45.0	631	841	23.5	10.60	1.720	22.9	32.3

Machine Description	Xmark 93	PLBwire9 3	PLBsurf9 3	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
44P-170 (400 MHz) AIX 4.3.3.1 L2 - 4MB	48.9	742	954	28.2	12.80	2.120	27.9	37.8
43P-260 (1-way) AIX 4.3.3.1 L2 - 4MB	34.9	433	674	14.8	6.46	0.975	13.1	24.3
43P-260 (2-way) AIX 4.3.3.1 L2 - 4MB	34.9	709	892	14.8	6.46	0.975	13.1	24.3
44P-270 (1-way) AIX 4.3.3.1 L2 - 4MB	47.7	728	936	27.1	11.90	1.960	26.2	33.7
44P-270 (2-way) AIX 4.3.3.1 L2 - 4MB	47.7	866	1140	27.1	11.90	1.960	26.2	33.7
44P-270 (4-way) AIX 4.3.3.1 L2 - 4MB	47.7	856	1170	27.1	11.90	1.960	26.2	33.7

10.1.2 POWER GXT2000P Graphics Accelerator (# 2823)

The GXT2000P fits into a 32-bit PCI bus for price/performance mid-range positioning in UNIX graphics workstation markets. The GXT2000P is full of features that assure the optimal transfer of visual information to the user. The features target the strong MCAD market where the low manufacturing cost of the chip and PC board will provide a very attractive price performing system.

The POWER GXT2000P Graphics Accelerator is a short PCI adapter card that attaches to one slot. It supports display resolutions up to 1920 x 1200 maximum and refresh rate up to 76 Hz.

10.1.2.1 Supported Machines

The following machines support the GXT2000P:

- 7043-140
- 7043-150
- 7044-170
- 7043-260
- 7043-270 (Upgrade from 7043-260)
- 7044-270

Table 114 lists the supported and most applicable display types and cables for the POWER GXT2000P.

Table 114. GXT2000P Display and Cable Table

Display Type	Cable Feature Code
5081	N/S
6091 - 16	N/S
6091 - 19	N/S
6091 - 19i	# 4217
6091 - 23	N/S
POWERdisplay 16	N/S
POWERdisplay 17	# 4217
POWERdisplay 19	# 4217
POWERdisplay 20	# 4217
P50	Cable included with display
P70	# 4238
P72	Cable included with display
P76	Cable included with display
P92	Cable included with display
G52	N/S
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
T55A White	N/S
T55A Black	N/S
9516-B03	# 4217
9516-B04	# 4217
9516-B23	# 4217
T85A White	Cable included with display

Display Type	Cable Feature Code
T85A Black	Cable included with display
Note: *N/S = This accelerator/display combination is not supported.	

The 24" Sony Model GDM-W900 Trinitron Wide Aspect Display is supported on the GXT2000P on the following RS/6000 systems: 43P-140, 43P-150, 43P-260, and 43P-270 (upgrade from 43P-260).

10.1.2.2 Performance Table

Table 115 and Table 116 provide the performance results of the POWER GXT2000P graphics accelerator.

Table 115. GXT2000P Performance Table - Historic Measures

Machine Description	Xmark 93	PLBwire 93	PLBsurf 93	CDRS- 04	DX- 04	DRV- 05	Light -02	Awadv -02	ProCDRS- 01
43P-140 (233 MHz) GXT3000P AIX 432 - 3/99 Update CD	33.01	146.2	304.7	57.87	6.97	3.84	0.75	8.77	6.30
43P-140 (332 MHz) GXT3000P AIX 432 - 3/99 Update CD	37.15	170.3	359.7	65.37	8.65	4.81	1.00	11.34	7.31
43P-150 (375 MHz) GXT3000P AIX 432 - 3/99 Update CD	43.10	258.8	477.2	103.70	11.34	6.37	1.21	14.22	11.13
43P-260 (200 MHz) GXT3000P AIX 432 - 3/99 Update CD	42.61	448.1	561.5	188.91	13.27	7.06	1.13	12.09	17.61
43P-260 (2X200 MHz) GXT3000P AIX 432 - 3/99 Update CD	42.61	731.7	801.9	188.91	13.27	7.06	1.13	12.09	17.61

Table 116. GXT2000P Performance Table - Current Measures

Machine Description	Xmark 93	PLBwire 93	PLBsurf 93	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
43P-140 (233 MHz) GXT2000P AIX 4.3.3.1	31.6	144	293	6.68	3.71	0.686	9.28	6.21
43P-140 (332 MHz) GXT2000P AIX 4.3.3.1	35.5	167	347	8.37	4.67	0.919	12.2	7.13
43P-150 (375 MHz) GXT2000P AIX 4.3.3.1	42.5	257	463	10.9	6.12	1.14	15.2	11.1

Machine Description	Xmark 93	PLBwire9 3	PLBsurf9 3	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
44P-170 (333 MHz) GXT2000P AIX 4.3.3.1	50.0	652	716	19.6	10.0	1.700	23.1	23.1
44P-170 (400 MHz) GXT2000P AIX 4.3.3.1	55.3	761	820	23.4	12.0	2.080	27.6	24.7
43P-260 (1-way) GXT2000P AIX 4.3.3.1	40.2	440	537	11.6	6.08	0.969	13.0	17.0
43P-260 (2-way) GXT2000P AIX 4.3.3.1	40.2	730	760	11.6	6.08	0.969	13.0	17.0
44P-270 (1-way) GXT2000P AIX 4.3.3.1	53.8	744	801	22.1	11.30	1.930	26.3	24.2
44P-270 (2-way) GXT2000P AIX 4.3.3.1	53.8	849	1060	22.1	11.30	1.930	26.3	24.2
44P-270 (4-way) GXT2000P AIX 4.3.3.1	53.8	838	1180	22.1	11.30	1.930	26.3	24.2

10.1.3 POWER GXT300P Graphics Accelerator (# 2841)

The POWER GXT300P Graphics Accelerator provides a new level of productivity to day-to-day 2D visualization. It is designed to be the graphics accelerator for 2D applications that require full 24-bit color, multiple hardware color tables, and speed, providing the power for demanding applications with winning price/performance and function.

When configured as a part of an RS/6000 43P Model 140, 150, 260 or 44P Model 170, 270 workstation, the POWER GXT300P Graphics Accelerator produces an outstanding combination of 2D application throughput and user interaction for such application areas as:

- Electronic computer-aided design
- Geographical information systems and mapping
- Weather and air traffic control
- Command and control
- Process control

The POWER GXT300P Graphics Accelerator is a short PCI adapter card requiring one PCI slot and supports display resolutions of up to 1600 x 1200 and refresh rates up to 85 Hz.

10.1.3.1 Supported Machines

The following machines support the GXT300P:

- 7043-140
- 7043-150
- 7044-170
- 7043-260
- 7043-270 (Upgrade from 7043-260)
- 7044-270

AIX Version 4.3.3, or later, is required.

Table 117 lists the supported and most applicable display types and cables for the POWER GXT300P.

Table 117. GXT300P Display and Cable Table

Display Type	Cable Feature Code
5081	N/S
6091 - 16	N/S
6091 - 19	N/S
6091 - 19i	N/S
6091 - 23	N/S
POWERdisplay 16	N/S
POWERdisplay 17	# 4217
POWERdisplay 19	N/S
POWERdisplay 20	# 4217
P50	N/S
P70	# 4238
P72	Cable included with display
P76	Cable included with display
P92	Cable included with display

Display Type	Cable Feature Code
G52	N/S
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
T55A White	N/S
T55A Black	N/S
9516-B03	# 4217
9516-B04	# 4217
T85A White	N/S
T85A Black	N/S
Note: *N/S = This accelerator/display combination is not supported.	

10.1.3.2 Performance Table

Table 118 provides the performance results of the POWER GXT300P graphics accelerator.

Table 118. GXT300P Performance Table

Machine Description	Xmark 93	PLBwire9 3	PLBsurf9 3	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
43P-140 (233 MHz) GXT300P AIX 4.3.3.1	32.4	34.0	27.6	1.10	0.765	0.148	1.65	0.714
43P-140 (332 MHz) GXT300P AIX 4.3.3.1	36.3	35.9	29.0	1.33	0.936	0.177	1.98	0.897
43P-150 (375 MHz) GXT300P AIX 4.3.3.1	42.3	51.2	44.0	1.76	1.250	0.241	2.65	1.120
44P-170 (333 MHz) GXT300P AIX 4.3.3.1	50.7	139.0	110.0	2.56	1.930	0.398	3.95	1.340
44P-170 (400 MHz) GXT300P AIX 4.3.3.1	55.3	165.0	137.0	3.18	2.350	0.496	4.84	1.610

Machine Description	Xmark 93	PLBwire9 3	PLBsurf9 3	DX-05	DRV-0 6	Light-0 3	AWadvs-0 3	ProCDRS-0 2
43P-260 (1-way) GXT300P AIX 4.3.3.1	40.4	101.0	83.4	1.60	1.200	0.240	2.38	0.821
44P-270 (1-way) GXT300P AIX 4.3.3.1	54.0	157.0	131.0	3.00	2.280	0.475	4.58	1.520

All GXT300P SPECplb and SPECopc results run with SoftGraphics.

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

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

The POWER GXT250P supports 8-bit color depth that allows simultaneous display of up to 256 colors from a palette of approximately 16.7 million colors.

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

10.1.4.1 Supported Machines

The following machine support the GXT250P and GXT255P:

- 7043-270 (Upgrade from 7043-260)

Table 119 provides the supported display types and cables for the POWER GXT250P and GXT255P.

Table 119. GXT250P and GXT255P Display and Cable Table

Display Type	Cable Feature Code
5081	# 4239
6091 - 16	# 4239
6091 - 19	# 4239
6091 - 19i	# 4239
6091 - 23	# 4239
POWERdisplay 16	# 4239
POWERdisplay 17	# 4239
POWERdisplay 19	# 4239
POWERdisplay 20	# 4239
P50	Cable included with display
P70	# 4238
P72	Cable included with display
P76	Cable included with display
P92	Cable included with display
G52	N/S
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
T55A White	Cable included with display
T55A Black	Cable included with display
9516-B03	# 4217
9516-B04	# 4217
9516-B23	# 4217
T85A White	N/S

Display Type	Cable Feature Code
T85A Black	N/S
Note: * N/S = This accelerator/display combination is not supported.	

10.1.5 POWER GXT130P Graphics Accelerator (# 2830)

The GXT130P accelerates X Windows graphics applications allowing you to display business graphics, database status, and activity for securities trading quicker. A network control display can monitor activity and assist in load balancing, performance analysis, and capacity planning.

The GXT130P is also ideal for process control. In selected configurations, up to four monitors can be operated concurrently from the same RS/6000 system to display graphics and text for dozens of process points. For decision support systems, the GXT130P offers rapid, simultaneous display of large amounts of critical information at high resolution on up to four screens.

It satisfies the ISO 9241 standard of visual quality, refreshing the screen at up to 85 Hz for all resolutions, including 1600 x 1200 8-bit pixels. The GXT130P graphics accelerator is a half-length PCI card with a standard 15-pin DIN output connector.

AIX Version 4.2.1, 4.3.2, or higher is required with the appropriate additional device drivers.

The GXT130P supports the following 8-bits per pixels screen resolutions of:

- 640 x 480
- 800 x 600
- 1024 x 768
- 1280 x 1024
- 1600 x 1200

10.1.5.1 Supported Machines

The following machines support the GXT130P:

- 7043-140
- 7043-150
- 7043-260
- 7043-270 (Upgrade from 7043-260)

- 7044-170
- 7044-270
- 7046-B50
- 7025-F50
- 7025-F80
- 7026-H50
- 7026-H70
- 7026-H80
- 7026-M80
- 7013-S7X
- 7015-S7X
- 7017-S7X
- 7017-S80

Table 120 lists the supported and most applicable display types and cables for the POWER GXT130P.

Table 120. GXT130P Display and Cable Table

Display Type	Cable Feature Code
5081	N/S
6091 - 16	N/S
6091 - 19	N/S
6091 - 19i	N/S
6091 - 23	N/S
POWERdisplay 16	N/S
POWERdisplay 17	# 4217
POWERdisplay 19	N/S
POWERdisplay 20	# 4217
P50	Cable included with display
P70	# 4238
P72	Cable included with display
P76	Cable included with display

Display Type	Cable Feature Code
P92	Cable included with display
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
9516-B03	# 4217
9516-B04	# 4217
9516-B23	# 4217
T55A White	Cable included with display
T55A Black	Cable included with display
T85A White	Cable included with display
T85A Black	Cable included with display
Note: *N/S = This accelerator/display combination is not supported.	

10.1.5.2 Performance Table

Table 121 provides the performance results of the POWER GXT130P graphics accelerator. Advanced performance benchmarks were not tested on this accelerator given its target role as an entry-level accelerator.

Table 121. GXT130P Performance Table

Machine description	Xmark93
43P-140 (233 MHz), GXT130P, AIX 4.3.3.1	12.9
43P-140 (332 MHz), GXT130P, AIX 4.3.3.1	14.2
43P-150 (375 MHz), GXT130P, AIX 4.3.3.1	15.4
44P-170 (333 MHz), GXT130P, AIX 4.3.3.1	11.6
44P-170 (400 MHz), GXT130P, AIX 4.3.3.1	12.3
43P-260 (1-way 200 MHz), GXT130P, AIX 4.3.3.1	11.1
43P-260 (2-way 200 MHz), GXT130P, AIX 4.3.3.1	11.1
44P-270 (1-way 375 MHz), GXT130P, AIX 4.3.3.1	12.0

Machine description	Xmark93
44P-270 (2-way 375 MHz), GXT130P, AIX 4.3.3.1	12.0
44P-270 (4-way 375 MHz), GXT130P, AIX 4.3.3.1	12.0

10.1.6 POWER GXT120P Graphics Accelerator (# 2838)

The POWER GXT120P is a versatile, low-priced graphics accelerator for RS/6000 PCI-bus 43P/44P Models 140, 150, 240, 260, and 270 (upgrade from 260), and for Models F50, H50, H70, H80 (upgrade from H70), S70, S70 Advanced, and S80. It accelerates X Windows graphics applications, such as displaying business graphics, database 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 used also 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.

10.1.6.1 Supported Machines

The following machines support the GXT120P:

- 7043-140
- 7043-150
- 7043-260
- 7043-270 (Upgrade from 7043-260)
- 7025-F50
- 7026-H50
- 7026-H70
- 7026-H80 (Upgrade from 7026-H70)
- 7013-S7X
- 7015-S7X
- 7017-S7X

Table 122 lists the supported and most applicable display types and cables for the POWER GXT120P.

Table 122. GXT120P Display and Cable Table

Display Type	Cable Feature Code
5081	N/S
6091 - 16	N/S
6091 - 19	N/S
6091 - 19i	N/S
6091 - 23	N/S
POWERdisplay 16	N/S
POWERdisplay 17	# 4217
POWERdisplay 19	N/S
POWERdisplay 20	# 4217
P50	Cable included with display
P70	# 4238
P72	Cable included with display
P76	Cable included with display
P92	Cable included with display
G52	Cable included with display
P200	# 4238
P201	# 4237
P202	Cable included with display
P260	Cable included with display
T55A White	Cable included with display
T55A Black	Cable included with display
9516-B03	# 4217
9516-B04	# 4217
9516-B23	# 4217
T85A White	Cable included with display

Display Type	Cable Feature Code
T85A Black	Cable included with display
Note: *N/S = This accelerator/display combination is not supported.	

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

The Ultimedia Video Capture Adapter/S (# 2639) is a PCI-based adapter that allows the capture of live video and still images. It supports three industry-standard video formats:

- NTSC
- PAL
- SECAM

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

10.1.7.1 Supported Machines

The following machines support the PCI Ultimedia Video Capture Adapter:

- 7043-140
- 7043-150
- 7044-170
- 7043-260
- 7043-270 (Upgrade from 7043-260)
- 7044-270

10.1.8 Percentage Performance Improvement

Table 123 shows the performance improvement in percentages, comparing the 7044-270 with the 7043-260 and the 7044-170 with the 7043-150 systems, configured with GXT3000P or GXT2000P graphics accelerators.

Table 123. Performance Improvement in Percentages

	7044-270 vs. 7043-260 1-Way	7044-270 vs. 7043-260 1-Way	7044-170 (333) vs. 7043-150	7044-170 (333) vs. 7043-150	7044-170 (400) vs. 7043-150	7044-170 (400) vs. 7043-150
	2000P	3000P	2000P	3000P	2000P	3000P

	7044-270 vs. 7043-260 1-Way	7044-270 vs. 7043-260 1-Way	7044-170 (333) vs. 7043-150	7044-170 (333) vs. 7043-150	7044-170 (400) vs. 7043-150	7044-170 (400) vs. 7043-150
PLBsurf	49%	39%	55%	69%	77%	92%
PLBwire	69%	68%	154%	148%	196%	192%
Awadvs-06	102%	100%	52%	46%	82%	78%
DRV-06	86%	84%	63%	73%	96%	108%
DX-05	98%	78%	80%	124%	115%	169%
Light-03	99%	101%	49%	51%	82%	86%
ProCDRS-02	42%	45%	108%	183%	123%	204%
x11perf	34%	37%	18%	24%	30%	35%
Average	69%	67%	68%	83%	95%	113%

Analysis of the data in Table 123 concludes the 375 MHz POWER3-II processor has approximately 67 percent better performance than the 200 MHz POWER3 using GXT2000P or GXT3000 graphics accelerators. The 7044-170 333 MHz system is 68 percent better than the 7043-150 375 MHz system using a GXT2000P and 83 percent better with GXT3000P. The 400 MHz Model 170 is 95 percent better than the 375 MHz Model 150 using a GXT2000P and 113 percent better with a GXT3000P.

10.2 Withdrawn Graphics Accelerators

Table 124 provides a list of the graphics accelerators that are withdrawn. The withdrawn dates are for the US.

Table 124. Non-Supported Graphics Accelerators

Description	Feature Code	Withdrawn Date
POWER Gt3	# 2777	April 30, 1993
High-Performance 8-Bit 3D Color Graphics Processor	# 2780	May 28, 1993
High-Performance 24-Bit 3D Color Graphics Processor	# 2781	May 28, 1993
POWER Gt4x 8-Bit	# 2790	December 21, 1993

Description	Feature Code	Withdrawn Date
POWER Gt4x 24-Bit	# 2791	December 21, 1993
POWER Gt4 8-Bit	# 2795	December 21, 1993
POWER Gt4 24-Bit	# 2796	December 21, 1993
Grayscale Graphics Display Adapter	# 2760	June 17, 1994
POWER GTO Accelerator	# 4350	November 4, 1994
Color Graphics Display Adapter	# 2770	January 6, 1995
POWER Gt1x	# 4207	September 19, 1995
POWER Gt1	# 4208	September 19, 1995
POWER Gt3i	# 2768	September 19, 1995
IBM E15-type Graphic	# 2731	January 19, 1996
POWER Gt4e	# 2776	October 25, 1996
S15 Graphics Accelerator	# 2657	January 21, 1997
POWER GXT155L	# 2665	July 18, 1997
POWER GXT150L	# 2660	September 24, 1997
POWER GXT500P	# 2854	January 30, 1998
POWER GXT1000	# 7252	March 31, 1998
POWER GXT110P	# 2839	August 14, 1998
POWER GXT1000	7250-001/002	July 24, 1998
POWER GXT800M	# 2850	March 19, 1999
MVP	# 2837	March 19, 1999
POWER GXT255P	# 2852	September 8, 1999
POWER GXT550P	# 2845	September 8, 1999
POWER GXT550P	# 2855	September 8, 1999
POWER GXT800P	# 2853	September 8, 1999
POWER GXT800P	# 2859	September 8, 1999
POWER GXT120P	# 2838	December 13, 1999
POWER GXT250P	# 2851	December 13, 1999

10.3 Limitations

Table 125 summarizes the maximum number of accelerators per system and some possible limitations.

Table 125. Maximal Number of Accelerators Per System and Limitations

System	130P	300P	2000P	3000P
7043P-140	4	4	4	N/A
7043P-150	4 ^{3,5}	4 ³	4	1 ²
7043P-260	4	4	4	2 ¹
7044P-170	4	4	4	2 ¹
7044P-270	4	4	4	2 ²
7046-B50	1	N/S	N/S	N/S
7025-F50	2 ⁵	N/S	N/S	N/S
7025-F80	1	N/S	N/S	N/S
7026-H50	2 ⁵	N/S	N/S	N/S
7026-H70	2 ⁵	N/S	N/S	N/S
7026-H80	1	N/S	N/S	N/S
7026-M80	1	N/S	N/S	N/S
7017-S7A	1 ⁴	N/S	N/S	N/S
7017-S80	1 ⁴	N/S	N/S	N/S

Note:
 * N/S = This accelerator is not supported.
¹ Must be located in slot 2, 4; if the POWER GXT3000P is installed in slot 4, slot 3 will be unavailable.
² Must be located in slot 3; the POWER GXT3000P takes the physical space of two adapters.
³ If both the POWER GXT120P or GXT130P and the POWER GXT255P or POWER GXT300P are installed, the POWER GXT255P or POWER GXT300P must be installed in slot 2 or 3.
⁴ It is strongly recommended you locate GXT130P accelerators in the primary drawer. Do not use it in I/O drawer 3 or 4.
⁵ Do not install a PCI Dual Channel Ultra2 SCSI Adapter on the same PCI bus as a POWER GXT130P.

Note

Refer to the *PCI Adapter Placement Reference*, SA38-0538, for important information regarding permissible and optimum slot placement of your PCI graphics accelerators.

You should also be aware of the interaction with other PCI adapters in the same bus with your graphics accelerators. A particularly important point is the performance degradation that may be seen by placing the 32-bit adapter in the same bus as the 64-bit GXT3000P accelerator.

Chapter 11. SCSI, SSA, and Fibre Channel Overview

Standards committees drive the adoption of interface protocol allowing any peripheral device following the standard to be used interchangeably. These protocols are written in firmware programming languages and are resident on the interface chip sets.

This chapter describes common standards used with the RS/6000 family.

11.1 SCSI

SCSI, Small Computer Systems Interface, is a bus-level interface and is widely used in mid- to high-performance workstations and servers. SCSI offers faster transfer rates than ATA/IDE, the interface most commonly used in desktop PCs.

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

Note

The American National Standards Institute (ANSI) refers to the different SCSI specifications using the SCSI-I, SCSI-II, and SCSI-III type nomenclature. IBM uses SCSI-1, SCSI-2, Ultra SCSI, Ultra2 SCSI, and Ultra3 SCSI nomenclature in official product names.

11.1.1 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.

11.1.2 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. Other enhancements are discussed in the text that follows. SCSI-II comes in many varieties, such as SCSI-II, SCSI-II Fast, and SCSI-II Fast/Wide.

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

11.1.2.1 Common Command Set

The SCSI-II standard defines a set of commands that must be interpreted by all devices that are attached to an 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.

11.1.2.2 Tagged Command Queuing

Tagged Command Queuing (TCQ) is an 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.

11.1.3 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:
 - Fibre Channel Physical and Signaling Interface (FC)
 - IEEE 1394 High Performance Serial Bus
 - 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)

SCSI-III is 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 (IEEE 1394). 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.

11.1.3.1 Overview of SCSI-III Standards

The breakdown of the SCSI-III standards is shown in Figure 67 on page 394 and is described in Table 126 on page 394.

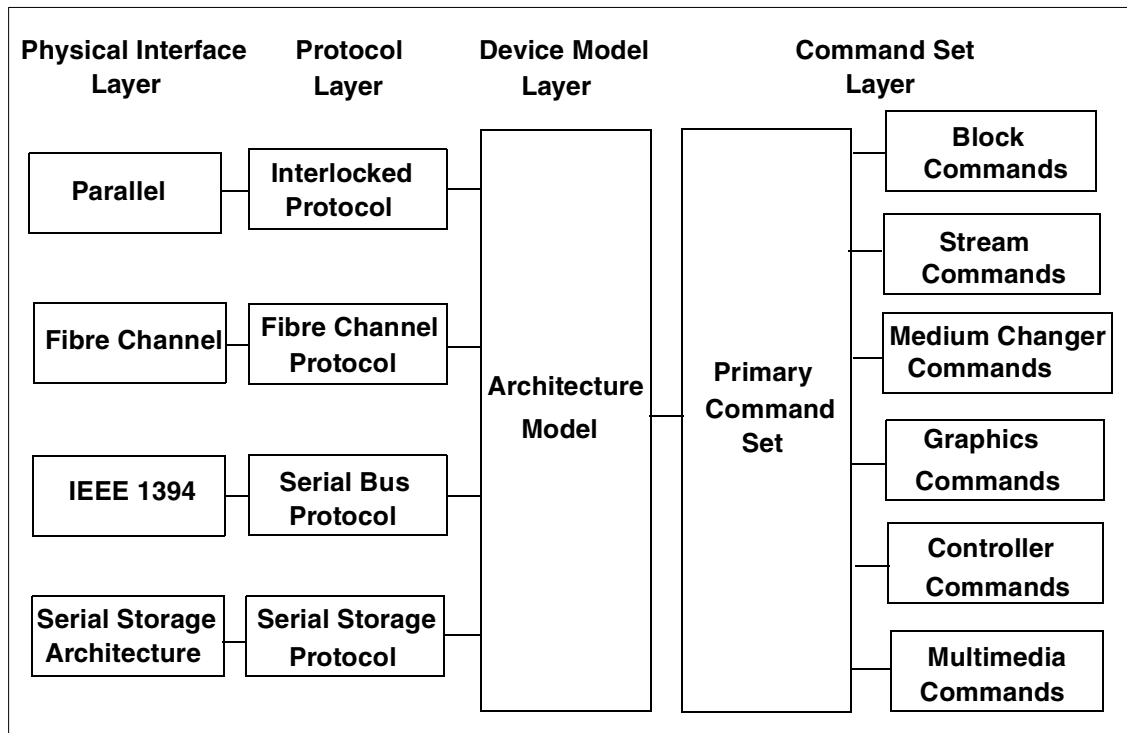


Figure 67. SCSI III Standards Overview

Table 126. Overview of SCSI-III Standards

Standards	Abbreviation	Function
SCSI-III Interlocked Protocol	SIP	Describes the protocols used on the SPI bus
Fibre Channel Protocol	FCP	Describes the protocols used on the FC bus
Serial Bus Protocol	SBP	Describes the protocols used on the IEEE 1394 bus
Serial Storage Protocol	SSP	Describes the protocols used on the SSA bus
SCSI-III Architecture Model	SAM	Describes the architecture of the SCSI-III model (includes the SCSI-III device models)
SCSI-III Primary Commands	SPC	Describes the commands that all SCSI devices must implement
SCSI-III Block Commands	SBC	Describes the commands used to transmit blocks of data

Standards	Abbreviation	Function
SCSI-III Stream Commands	SSC	Describes the commands used to transmit <i>streams</i> of data
SCSI-III Medium Changer Commands	SMC	Describes the commands used to change a medium in a device
SCSI-III Graphic Commands	SGC	Describes the commands that involve graphics
SCSI-III Controller Commands	SCC	Describes the commands used to configure and test the controller of a device
SCSI-III Multimedia Commands	MMC	Describes commands that involve multimedia data, such as audio and video

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

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

SPI is as close to SCSI-II technology as possible 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 from 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 backward compatible with SCSI-I and SCSI-II. It is ideal for those users that have sizable investments in SCSI-II equipment because SPI integrates SCSI-II and SCSI-III devices on the same chain. The use of SCSI-III provides much more function due to the improvements in the command set and data transfer rates. The migration from SCSI-II to SPI is relatively simple because all the connectors, cables, and terminators essentially remain the same.

The SCSI-III Parallel Interface specification uses the terms *Fast-20 Narrow* and *Fast-20 Wide*. Fast-20 Narrow is an 8-bit bus with a maximum data transfer rate of 20 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 the bus, 20 MHz, which is double the speed of the SCSI-II Fast.

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

After SCSI, there was Ultra SCSI, then Ultra2 SCSI, or Low Voltage Differential (LVD), a disk drive interface that is faster and more reliable than previous SCSI standards. Ultra2 SCSI provides the significant advantages of increased bus bandwidth, faster transfer rates, greater device connectivity, and better configuration flexibility. Best of all, Ultra2 SCSI is backward compatible with earlier SCSI versions.

Now, an even more powerful interface is available: Ultra3 SCSI. As ratified by the SCSI Trade Association (STA), Ultra3 SCSI refers to products that incorporate any or all of the following features of the SCSI SPI-3 standards:

- Double transition clocking
- Domain validation
- Cyclic redundancy check (CRC)
- Packetization
- Quick arbitration select (QAS)

A specific feature set of Ultra3 SCSI, known throughout the industry as Ultra160 SCSI, is currently taking SCSI to new levels of performance. Named for its superior 160 MBps data transfer speed, the latest generation of SCSI technology incorporates the three management features of Ultra3 SCSI that specifically affect data transfer: CRC, domain validation, and double transition clocking.

These new capabilities provide a cost-effective way to dramatically boost both device performance and reliability. Because Ultra160 SCSI is compatible with Ultra2 SCSI devices, it helps protect existing investments and ensures a smoother transition. The cables, connectors, and terminators are the same for both Ultra160 SCSI and Ultra2 SCSI. In fact, Ultra160 SCSI host controllers can support Ultra2 SCSI devices, which enables the mixing of Ultra160 and Ultra2 SCSI devices on the same bus. When Ultra160 SCSI and Ultra2 SCSI devices are mixed, each device can at operate at its full rated speed.

A Wide Variety of Benefits

By incorporating a specific combination of Ultra3 SCSI optimized capabilities, Ultra160 SCSI has established an effective industry standard for performance and device compatibility. The Ultra160 SCSI feature set has already been widely embraced by leading system manufacturers and SCSI drive and component suppliers, such as Adaptec, Compaq, Fujitsu, and Hewlett Packard. As one of the leading companies in the storage industry, IBM is building on its previous commitment to SCSI technology by incorporating Ultra160 SCSI into its products, such as the IBM Ultrastar hard disk drive family.

Note

Currently, the RS/6000 family does not support Ultra160 SCSI.

11.1.3.2 General Terminology

The following is a list of general terminology for SCSI:

- | | |
|-------------------|--|
| Fast | Fast refers to the doubling of the data transfer rate from the SCSI 5 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 <i>burst</i> 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 <i>Fast-20 Narrow</i> . <i>Fast-20 Wide</i> is the 16-bit version, also known as Ultra SCSI. |
| Ultra SCSI | Ultra SCSI, as described above, is a subset of the SCSI-III specification. It is effectively a Fast SCSI bus running at 20 MHz. Ultra SCSI can produce a maximum transfer of 20 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 Ultra2 is the latest extension to the SCSI-III specification that 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).

Ultra3 SCSI Ultra3 SCSI - known throughout the industry as Ultra160 SCSI - is currently taking SCSI to new levels of performance. Named for its superior 160 MBps data transfer speed, the latest generation of SCSI technology incorporates the three management features of Ultra3 SCSI that specifically affect data transfer: CRC, domain validation, and double transition clocking.

Note

Wide refers to the width of the bus between the SCSI Adapter and its attached devices. Do not confuse this with the width of the host bus interface.

11.1.3.3 Asynchronous versus Synchronous

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

11.1.3.4 SCSI Differential

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

Consider Figure 68 on page 399 as an example. Signal A is transmitted along two wires, A+ and A-. A- is the mirror image of A+. When some line noise appears, such as a signal spike, the spike will be of the same polarity on both

A+ and A-. On the receiving end, the signals are subtracted, and the result is A+/A- as shown in the diagram. The spike is canceled out by the subtraction.




Figure 68. SCSI Signal Comparison

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

11.1.4 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 127 on page 400, the maximum length of the SCSI bus for a SCSI-II F/W configuration is 3.0 meters. With a repeater, it is possible to extend this distance.

11.1.5 Summary of SCSI Specifications

SCSI defines many different modes of operation, including several different data transfer modes. The most common are shown in Table 127.

Table 127. SCSI Specification Summary

	SCSI	Fast SCSI	Ultra SCSI	Wide Ultra SCSI	Ultra2 SCSI	Wide Ultra2 SCSI	Ultra3 SCSI
Data Transfer Rates (MBps)	5	10	20	40	40	80	160
Max. Data Bus Width (bits)	8-bit	8-bit	8-bit	16-bit	8-bit	16-bit	16-bit
Max. Cable Length (meters)	6	3	1.5 - 3	1.5 - 3	12	12	12
Max. Device Support	8	8	8 - 4	8 - 4	8	16	16
Note: * These are industry standard interface names. Other companies may use different names for the same interfaces mentioned here. Take note that the technical data is the distinguishing factor.							

11.1.6 General SCSI Cabling Considerations

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

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

11.1.6.1 Default SCSI Cable Routing

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

11.1.6.2 SCSI Bus Length

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

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

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

Devices that have two SCSI connectors have internal cabling that must be included when calculating the total cable length. When attaching these devices to a SCSI chain, connect one cable to one connector and the other cable to the other connector. Do not piggy back the second cable/terminator onto the first as you would with a device that had only one SCSI connector as shown in Figure 69.

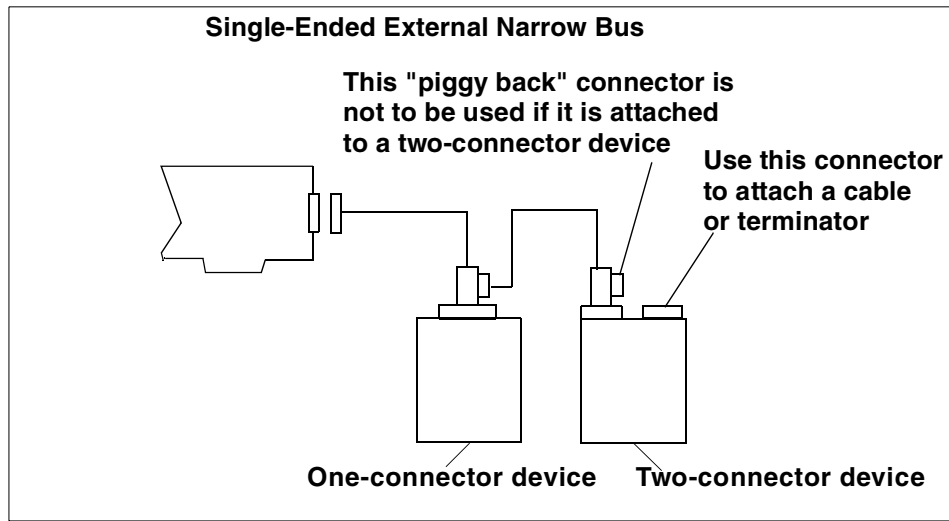


Figure 69. Piggy-Back Connectors

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

11.1.6.3 SCSI Terminators

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

- For SCSI Adapters with external devices only, make sure the appropriate SCSI terminator is connected to the last device on the bus.
- For SCSI Adapters with internal devices only, make sure that the appropriate SCSI terminator is connected to the end of the internal SCSI cable.
- For SCSI Adapters 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.

11.1.6.4 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.

Some examples of cabling the PCI Single-Ended Ultra SCSI Adapter (# 6206) and PCI Differential-Ended Ultra SCSI Adapter (# 6207) you can find in Appendix B, "SCSI Cabling Examples" on page 633.

11.2 Serial Storage Architecture Overview

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

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

- Enables hot-swapping and automatic configuration of new devices (without interrupting file service to users).
- Supports up to 128 total devices (including multiple initiator) in a single SSA configuration, compared to a maximum of 16 devices for SCSI.
- Provides configuration flexibility with up to 25 meters of separation between devices.
- Maximizes data transfer, currently enabling 160 MBps of available bandwidth per loop.
- Distances of up to 10 km are possible with a 7133-D40 or 7133-T40 and an optical extender.

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.

11.3 Comparison between SCSI and SSA Architecture

Table 128 lists the common features found in SCSI and SSA devices for comparison, and Figure 70 shows how the subsystems differ.

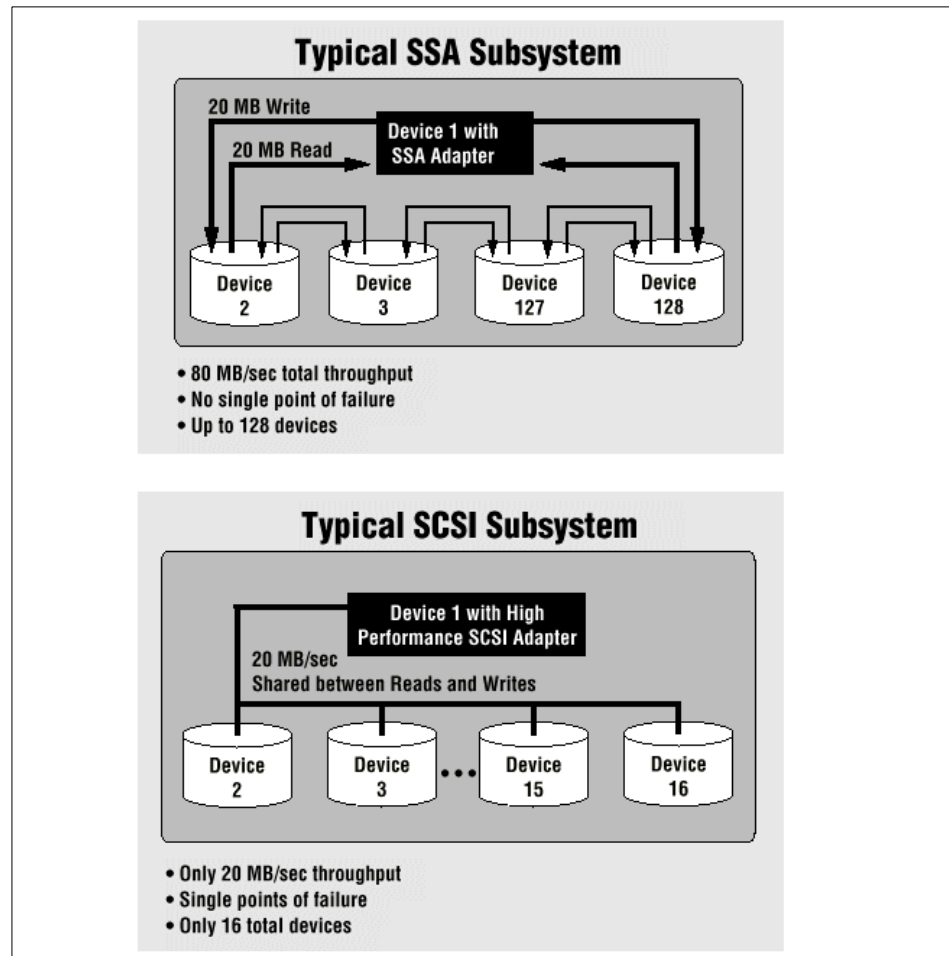


Figure 70. Comparison between SCSI and SSA Subsystems

Table 128. Comparing SSA and SCSI

Feature	SCSI	SSA
Cable Type	68 conductor shielded cable (SCSI-II Differential Fast/Wide)	4 conductor twisted-pair

Feature	SCSI	SSA
Maximum Cable Lengths	25 meters total	25 meters between devices, up to 2500 meters between devices with the use of a Fiber Optic Extender
Cable Costs	Significant	Low
Individual Drive Addressing	Must be set manually	Dynamic
Cable Termination	Electrical terminator must be at the end of the SCSI bus	No termination required, closed serial loops
Fault Tolerance	If single data path (bus) is compromised, access is lost	No single point of failure. Multiple data paths prevents a single loop from causing loss of access to information
Maximum Throughput	Up to 80 MBps	Up to 160 MBps

11.4 Fiber Channel-Arbitrated Loop (FC-AL)

Fibre Channel-Arbitrated Loop (FC-AL) is an exceptionally high-bandwidth industry-standard interface primarily targeted toward high-end servers and similar demanding applications. FC-AL uses fiber optic cabling in a loop configuration to produce maximum transfer speeds of 100 MBps and is designed to connect up to 127 devices as far as 10 kilometers apart, enabling data storage in remote, secure locations distant from the server.

FC-AL devices can be dual ported, providing two simultaneous input/output sessions that doubles maximum throughput, and FC-AL enables *hot swapping*, so you can add and remove hard drives without interrupting system operation, an important option in server environments. FC-AL adapters tend to cost more than SCSI or SSA adapters.

Chapter 12. Internal Storage Devices

This chapter provides an overview of the internal SCSI devices for the available RS/6000s and their main characteristics:

- Disk drives
- CD-ROM drives
- Tape drives

12.1 Disk Drives

Table 129 lists the available disk drives that are supported on RS/6000 systems.

Table 129. Available Disk Drives and Feature Codes

Feature Code	Disk Drive
2502	9.1 GB Ultra SCSI 1" High (25 mm) Disk Drive
2503	18.2 GB Ultra SCSI 1" High (25 mm) Disk Drive
2906	9.1 GB 10,000 RPM Ultra SCSI 16-bit Enhanced Disk Drive
2908	9.1 GB Ultra SCSI 16-Bit 1" High (25 mm) Disk Drive
2909	18.2 GB Ultra SCSI 1" High (25 mm) Disk Drive
2913	9.1 GB Ultra SCSI 16-Bit Hot Swap Disk Drive
3002	9.1 GB 10,000 RPM Ultra SCSI 1" High Hot Swap Disk Drive
3021	9.1 GB Ultra2 SCSI 16-Bit 1" High (25 mm) Disk Drive Assembly
3023	18.2 GB Ultra2 SCSI 16-Bit 1" High (25 mm) Disk Drive Assembly
3025	9.1 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly
3026	18.2 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly
3027	9.1 GB 10,000 RPM Ultra SCSI Disk Drive
3029	9.1 GB Ultra SCSI 16-Bit Enhanced Disk Drive
3102	18.2 GB 10,000 RPM Ultra SCSI Disk Drive
3103	18.2 GB Ultra SCSI Enhanced Disk Drive
3104	18.2 GB Ultra SCSI 1" High (25 mm) Hot-Swap Disk Drive

Feature Code	Disk Drive
3110	18.2 GB 10,000 RPM Ultra SCSI 16-Bit Enhanced Disk Drive
3117	18.2 GB 10,000 RPM Ultra SCSI Hot-Swap Disk Drive
3119	36.4 GB 10,000 RPM Ultra SCSI Disk Drive
3128	36.4 GB 10,000 RPM Ultra SCSI 16-bit Enhanced Disk Drive
3129	36.4 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly
3070	¹ 9.1 GB Serial Storage Architecture 10K RPM Hot Swap Disk Drive
3077	18.2 GB Serial Storage Architecture 10K RPM Disk Drive Assembly
3079	9.1 GB Serial Storage Architecture 10K RPM Disk Drive Assembly
Note: ¹ Check for latest availability information	

12.1.1 9.1 GB Ultra-SCSI 1" High Disk Drive (# 2502)

The 9.1 GB Ultra-SCSI 1-inch (25mm) High Disk Drive offers high-performance that provides 9.1 GB of storage capacity and supports the industry Ultra-SCSI interface speed of 40 MB/s. Requires Ultra-SCSI adapter to operate at 40 MB/s.

12.1.1.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSI Spec X3T10/1071D
- Average Seek Time: 7.2 ms (based on four READS and one WRITE)
- Average Latency: 4.17 ms
- Rotational Speed: 7200 RPM
- Maximum Media Data Transfer Rate: 22.4 MB/s

12.1.2 18.2 GB Ultra-SCSI 1" High Disk (# 2503)

The 18.2 GB Ultra-SCSI 1-Inch High (25 mm) Disk Drive offers high-performance that provides 18.2 GB of storage capacity and supports the industry Ultra-SCSI interface speed of 40 MB/s. To operate at 40 MB/s, an Ultra SCSI Adapter (# 6206) is required.

12.1.2.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSI Spec X3T10/1071D
- Average Seek Time: 7.2 ms (based on four READS and one WRITE)
- Average Latency: 4.17 ms
- Rotational Speed: 7200 RPM
- Maximum Media Data Transfer Rate: 32 MB/s

12.1.3 9.1 GB 10K RPM Ultra-SCSI 16-Bit Enhanced Disk (# 2906)

The 9.1 GB 10,000 RPM Ultra SCSI 16-Bit 1-Inch High (25 mm) Disk Drive provides 9.1 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.3.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.5 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 30.5 MB/s

12.1.4 9.1 GB Ultra-SCSI 16-Bit 1" High Disk (# 2908)

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

12.1.4.1 Characteristics

- Form Factor: 1 inch (25mm) high by 4.0 inch (102mm) wide
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)

- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 6.7 ms
- Average Latency: 4.17 ms
- Rotational Speed: 7200 RPM
- Maximum Media Data Transfer Rate: 22.4MB/s

12.1.5 18.2 GB Ultra-SCSI 1" High Disk (# 2909)

The 18.2 GB Ultra-SCSI 16-Bit 1 inch High (25 mm) Disk Drive provides 18.2 GB of storage capacity and supports the industry's Ultra-SCSI interface speed of 40 MB/s. An Ultra SCSI Adapter (# 6206) is required to operate at 40 MB/s.

12.1.5.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0-inch (102 mm) wide
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 6.7 ms
- Average Latency: 4.17 ms
- Rotational Speed: 7200 rpm
- Maximum Media Data Transfer Rate: 22.4 MB/s

12.1.6 9.1 GB Ultra-SCSI 16-Bit Hot Swap Disk Drive (# 2913)

The 9.1 GB Ultra-SCSI 16-Bit 1 inch High (25mm) Disk Drive offers high-performance that provides 9.1 GB of storage capacity and supports the industry Ultra-SCSI interface speed of 40 MB/s.

12.1.6.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 6.7 ms
- Average Latency: 4.17 ms
- Rotational Speed: 7200 RPM

- Maximum Media Data Transfer Rate: 22.4 MB/s

12.1.7 9.1 GB 10K RPM Ultra SCSI 1" High Hot Swap Disk (# 3002)

The 9.1 GB 10,000 RPM Ultra SCSI 16-bit 1 inch (25 mm) high Disk Drive provides 9.1 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.7.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.1 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.8 9.1 GB Ultra2 SCSI 16-Bit 1" High Disk (# 3021)

The 9.1 GB Ultra2 SCSI 16-bit 1 inch (25 mm) High Disk Drive Assembly provides 9.1 GB of storage capacity and supports the industry's Ultra SCSI interface speed of 80 MB/s. An Ultra2 SCSI adapter is required to operate at 80 MB/s.

12.1.8.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra2 SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 7.2 ms (based on four READS to one WRITE)
- Average Latency: 4.17 ms
- Rotational Speed: 7200 rpm
- Maximum Media Data Transfer Rate: 32 MB/s

12.1.9 18.2 GB Ultra2 SCSI 16-Bit High Disk Assembly (# 3023)

The 18.2 GB Ultra2 SCSI 16-bit 1 inch (25 mm) High Disk Drive Assembly provides 18.2 GB of storage capacity and supports the industry's Ultra SCSI

interface speed of 80 MB/s. An Ultra2 SCSI adapter is required to operate at 80 MB/s.

12.1.9.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 7.2 ms (based on four reads to one write)
- Average Latency: 4.17 ms
- Rotational Speed: 7200 rpm
- Maximum Media Data Transfer Rate: 32 MB

12.1.10 9.1 GB 10K RPM Ultra2 SCSI Disk Drive Assembly (# 3025)

The 9.1 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly provides 9.1 GB of storage capacity and supports the industry standard Ultra2 SCSI interface speed of 80 MB/s.

12.1.10.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra2 SCSI (16-bit, Low Voltage Differential)
- Attachment Industry Spec: SCSI-3 Fast 40
- Average Seek Time: 5.1 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.11 18.2 GB 10K RPM Ultra2 SCSI Disk Assembly (# 3026)

The 18.2 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly provides 18.2 GB of storage capacity and supports the industry standard Ultra2 SCSI interface speed of 80 MB/s.

12.1.11.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No

- External Interface: Ultra2 SCSI (16-bit, Low Voltage Differential)
- Attachment Industry Spec: SCSI-3 Fast 40
- Average Seek Time: 5.1 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.12 9.1 GB 10K RPM Ultra-SCSI Disk Drive (# 3027)

The 9.1 GB 10,000 RPM Ultra SCSI Disk Drive provides 9.1 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.12.1 Characteristics

- Form Factor: 3.5-inch, 1-inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.5 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 30.5 MB/s

12.1.13 9.1 GB Ultra-SCSI 16-bit Enhanced Disk Drive (# 3029)

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

12.1.13.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSI Spec X3T10/1071D
- Average Seek Time: 6.7 ms
- Average Latency: 4.17 ms

- Rotational Speed: 7200 RPM
- Maximum Media Data Transfer Rate: 22.4 MB/s

12.1.14 18.2 GB 10K RPM Ultra SCSI Disk Drive (# 3102)

The 18.2 GB 10,000 RPM Ultra SCSI Disk Drive provides 18.2 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.14.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.1 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.15 18.2 GB Ultra-SCSI Enhanced Disk Drive (# 3103)

The 18.2 GB Ultra-SCSI 16-bit Enhanced High Disk Drive provides 18.2 GB of storage capacity and supports the industry's Ultra-SCSI interface speed of 40 MB/s.

12.1.15.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra-SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 6.7 ms
- Average Latency: 4.17 ms
- Rotational Speed: 7200 rpm
- Maximum Media Data Transfer Rate: 22.4 MB/s

12.1.16 18.2 GB Ultra-SCSI 1" High Hot-Swap Disk (# 3104)

The 18.2 GB Ultra SCSI 16-Bit 1 inch High (25 mm) Hot Swap Disk Drive provides 18.2 GB of storage capacity and supports the industry's Ultra SCSI interface speed of 40 MB/s.

12.1.16.1 Characteristics

- Form Factor: 1 inch (25 mm) high by 4.0 inch (102 mm) wide
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 Fast 20/ANSISpec X3T10/1071D
- Average Seek Time: 7.2 ms (based on four READS to one WRITE)
- Average Latency: 4.17 ms
- Rotational Speed: 7200 rpm
- Maximum Media Data Transfer Rate: 32 MB/s

12.1.17 18.2 GB 10K RPM Ultra SCSI Enhanced Drive (# 3110)

The 18.2 GB 10,000 RPM Ultra SCSI 16-Bit 1 Inch High (25 mm) Disk Drive provides 18.2 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.17.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.1 ms
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.18 18.2 GB 10K RPM Ultra SCSI Hot Swap Disk Drive (# 3117)

The 18.2 GB 10,000 RPM Ultra SCSI 16-bit 1-inch (25mm) high Disk Drive provides 18.2 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.18.1 Characteristics

- Form Factor: 3.5-inch, 1-inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.1 ms (based on four READS to one WRITE)
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.19 36.4 GB 10K RPM Ultra SCSI Disk (# 3119)

The 36.4 GB 10,000 RPM Ultra SCSI Disk Drive provides 36.4 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.19.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D
- Average Seek Time: 5.1 ms (based on four READS to one WRITE)
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.20 36.4 GB 10K RPM Ultra SCSI 16-Bit Enhanced Disk (# 3128)

The 36.4 GB 10,000 RPM Ultra SCSI 16-Bit Enhanced Disk Drive provides 36.4 GB of storage capacity and supports the industry standard Ultra SCSI interface speed of 40 MB/s.

12.1.20.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra SCSI (16-bit, Single Ended)
- Attachment Industry Spec: SCSI-3 fast 20, ANSI Spec X3T10/1071D

- Average Seek Time: 5.1 ms (based on four READS to one WRITE)
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.21 36.4 GB 10K RPM Ultra2 SCSI Disk Assembly (# 3129)

The 36.4 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly provides 36.4 GB of storage capacity and supports the industry standard Ultra2 SCSI interface speed of 80 MB/s.

12.1.21.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Cable included: No
- External Interface: Ultra2 SCSI (16-bit, Low Voltage Differential)
- Attachment Industry Spec: SCSI-3 Fast 40
- Average Seek Time: 5.1 ms (based on four READS to one WRITE)
- Average Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Data Transfer Rate: 44.3 MB/s

12.1.22 9.1 GB SSA 10K RPM Hot Swap Disk (# 3070)

This feature supplies one 9.1 GB Serial Storage Architecture (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, Check Light Emitting Diodes (LED) that show the status of the drive.

12.1.22.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Formatted Capacity: 9.1 GB
- Average Seek Time: 5.1 ms (based on four READS and one WRITE)
- Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Media Data Transfer Rate: 44.3 MB/s
- SSA Data Transfer Rate: 20 MB/s Full Duplex

12.1.23 18.2 GB SSA 10K RPM Disk Assembly (# 3077)

This feature supplies one 18.2 GB Serial Storage Architecture (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, Check Light Emitting Diodes (LED) that show the status of the drive.

12.1.23.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Formatted Capacity: 18.2 GB
- Average Seek Time: 5.1 ms (based on four READS to one WRITE)
- Latency: 2.99 ms
- Rotational Speed 10,000 RPM
- Maximum Media Data Transfer Rate: 44.3 MB/s
- SSA Maximum Data Transfer Rate: 40 MB/s Full Duplex

12.1.24 9.1 GB Serial Storage Architecture 10K RPM Disk (# 3079)

This feature supplies one 9.1 GB Serial Storage Architecture (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, Check Light Emitting Diodes (LED) that show the status of the drive.

12.1.24.1 Characteristics

- Form Factor: 3.5 inch, 1 inch (25 mm) high
- Formatted Capacity: 9.1 GB
- Average Seek Time: 5.5 ms
- Latency: 2.99 ms
- Rotational Speed: 10,000 RPM
- Maximum Media Data Transfer Rate: 30.5 MB/s
- SSA Maximum Data Transfer Rate: 40 MB/s Full Duplex

12.2 CD-ROM Drive

The SCSI-II Internal CD-ROM drive is a tray-loading CD-ROM drive providing up to 4800 KB/s maximum media data transfer rate. It is a 5.25 inch half-high form factor, single-ended, 8-bit, multi-session capable, CD-ROM drive that

provides state-of-the-art performance and supports existing 600 MB CD-ROM discs. It will operate in either the vertical or horizontal positions. Table 130 summarized the CD-ROM drive characteristics.

Table 130. Internal 32X CD-ROM Drive

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

12.3 Tape Drives

Table 131 summarizes the internal tape drives that are currently available on the RS/6000.

Table 131. Internal 4 mm and 8 mm Tape Drives

Product Name	12/24 GB 4 mm Internal Tape Drive	20/40 GB 8 mm Internal Tape Drive	20/40 GB 16-bit 8 mm Internal Tape Drive (White Bezel)
Feature Code	# 6159	# 6156	# 6154
Tape Capacity (Native Mode)	12 GB	20 GB	20 GB
Tape Capacity (Compression Mode)	24 GB	40 GB	40 GB
Type of Cartridge	IBM 4 mm DDS-3 Data Cartridge	IBM 8 mm Data Cartridge	IBM 8 mm Data Cartridge
Data Transfer Rate	1.1 MB/s Native Mode, 2.2 MB/s Compression Mode	3 MB/s Native Mode, 6 MB/s Compression Mode	3 MB/s Native Mode, 6 MB/s Compression Mode
Interface	SCSI-II (Single-Ended)	SCSI-II 16-bit (Single-Ended)	SCSI-II 16-bit (Single-Ended)

Product Name	12/24 GB 4 mm Internal Tape Drive	20/40 GB 8 mm Internal Tape Drive	20/40 GB 16-bit 8 mm Internal Tape Drive (White Bezel)
Form Factor	5.25 Inch Half-High	5.25 Inch Half-High	5.25 Inch Half-High
Compatibility	2 GB Mode (R/W) 4 GB Compression (R/W) 4 GB Mode (R/W) 8 GB Compression (R/W) 12 GB Mode (R/W) 24 GB Compression (R/W)	2.3 GB Mode (R/O) 2.3 GB Mode Compression (R/O) 5 GB Mode (R/O) 5 GB Compression (R/O) 7 GB Mode (R/O) 7 GB Compression (R/O) 20 GB Mode (R/W) 20 GB Compression (R/W)	2.3 GB Mode (R/O) 2.3 GB Mode Compression (R/O) 5 GB Mode (R/O) 5 GB Compression (R/O) 7 GB Mode (R/O) 7 GB Compression (R/O) 20 GB Mode (R/W) 20 GB Compression (R/W)

Chapter 13. External Storage Architectures and Devices

This chapter describes the various IBM Open Systems Storage products that can be attached to RS/6000 machines. After an introduction to the products, the following types of external storage devices are discussed:

- External disk storage products
- Tape drive products
- Tape automation products
- Optical devices
- Storage management products

13.1 Introduction to IBM Open Systems Storage Products

IBMs Open Systems Storage products are distinguished by the Seascope and Serial Storage Architectures.

13.1.1 IBM Seascope: Storage Enterprise Architecture

IBM Seascope storage enterprise architecture is a blueprint for comprehensive storage solutions optimized for the connected world. The Seascope architecture outlines new concepts for storage by integrating leading technologies from IBM—including disk, tape, optical, RISC processors, and rich software functions—to provide highly-reliable, scalable, and versatile application-based storage solutions that span the range of servers from PCs to supercomputers.

The Seascope architecture is based on three major principles:

- An integrated storage server architecture, based on industry standards, to address the issues of complex data movement.
- *Snap-in* hardware and software building blocks that are key to keeping pace with advancements in function and technology as well as providing superior investment protection.
- A foundation for universal data access with a focus on overcoming the boundaries and complexities of accessing data in a connected world.

Seascope provides the strategic foundation for true data sharing, the ultimate goal for SAN technology.

13.1.2 Storage Area Network

A storage area network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

SANs are based on a *fabric* of Fibre Channel hubs, switches, and gateways connecting storage devices, such as disk arrays, optical disks or tape libraries, and servers on a many-to-many basis. Application and transaction servers are attached to both the SAN and to Local Area Networks (LANs) or Wide Area Networks (WANs), thus creating what appears to be a massive pool of data.

The move to storage area networks has been motivated by the need to manage the dramatically increasing volume of business data and to mitigate its effect on network performance. Key factors include:

- e-business -- Securely transforming internal business processes and improving business relationships to expedite the buying and selling of goods, services, and information through the Internet.
- Globalization -- The extension of IT systems across international boundaries.
- Zero latency -- The need to exchange information immediately for competitive advantage.
- Transformation -- The ability to continually adapt, while immediately accessing and processing information to drive successful business decisions.

Distributed computing, client/server applications, and open systems give today's enterprises the power to fully integrate hardware and software from different vendors to create systems tailored to their specific needs. These systems can be fast, efficient, and capable of providing a competitive edge.

Storage area networks remove data traffic, such as backup processes, from the production network, therefore, giving IT managers a strategic way to improve system performance and application availability.

Storage area networks improve data access. Using Fibre Channel connections, SANs provide the high-speed network communications and distance needed by remote workstations and servers to easily access shared data storage pools.

With a SAN, virtually unlimited expansion is possible with hubs and switches. Nodes can be removed or added with minimal disruption to the network.

13.1.3 RS/6000 and Storage Area Networks

IBM currently offers a wide range of SAN-ready RS/6000 servers, from entry desktop platforms, such as the F50, F80 and the rack-mounted H50, H70, and H80, through the S family of high-end Enterprise Servers and Large Scale SP systems.

The family of SAN-ready RS/6000 servers supports Fibre Channel attachment of the IBM Fibre Channel RAID Storage Server (except for the SP). This disk subsystem provides a superb solution for customers with information storage requirements of up to one terabyte. Recent enhancements include support of this subsystem by HACMP, IBM's industry-leading high availability software, in two-way configurations.

Support for Fibre Channel attachment of IBM Enterprise Storage Server (ESS) to RS/6000 SAN-ready servers is possible using the IBM SAN Data Gateway. ESS is a robust, highly-available, enterprise-wide storage solution that can scale up to 11 terabytes.

In addition, all RS/6000 SAN-ready servers support Fibre Channel attachment, also via the SAN Data Gateway, of the Magstar family of tape subsystems - the 3570, 3575 tape library, 3590, and 3494 tape library. This attachment also supports an STK silo if it utilizes a 3590 as its tape subsystem.

When attachment of supported subsystems using the SAN Data Gateway is implemented in conjunction with a pair of IBM Fibre Channel Storage Hubs, distances of up to 10 km can be attained, substantially enhancing the range of physical planning options.

13.1.4 Serial Storage Architecture

Serial Storage Architecture (SSA) is an industry standard specifically designed to meet the interface requirements of disk drives. SSA provides exceptional performance, flexibility, and high data availability by using a superior loop structure. Dual data paths to each device provide leading edge data access flexibility. SSA devices are easy to install and configure. SSA *hot-swappable* drives can be removed and replaced without disrupting communication between the adapter and other disks on the loop.

SSA offers increased data availability and overcomes numerous SCSI limitations, including number of devices, distance, and addressing.

Unlike SCSI bus configurations, SSA devices are configured in loops and do not require bus arbitration. This enables multiple concurrent operations to

occur in separate sections of the loop, therefore, resulting in higher overall throughput. This exclusive capability of SSA is referred to as *Spatial Reuse*.

13.2 IBM External Disk Storage

This section begins with a section on product positioning. Table 132 on page 425 summarizes the following IBM Open Systems Storage External Disk Systems:

- 2102 IBM Fibre Channel RAID Storage Server
- 7133 Serial Disk Systems
- Enterprise Storage Server
- 2104 Expandable Storage Plus

Following this are more detailed descriptions of the following IBM external disk storage products:

- 2102 IBM Fibre Channel RAID Storage Server
- 7133 Serial Disk System Advanced Models T40 and D40
- Enterprise Storage Server
- 2104 Expandable Storage Plus
- 7204 External SCSI Disk Drives
- 7203 Model 001 Portable Disk Drive

13.2.1 Disk Product Positioning

There is a full range of disk storage products for RS/6000 systems. For Internet Service Providers, Application Service Providers, or for customers who require a low-cost or entry-level storage solution, the IBM 2104 Expandable Storage Plus is available. It provides a compact, affordable SCSI storage drawer or tower for RS/6000 systems.

The IBM 7133 Serial Disk System Advanced Models D40 and T40 are storage systems based on third generation Serial Storage Architecture. They provide a great solution for high performance, high availability, and connectivity, with capacity that starts at 36 GB and can grow to many terabytes. The Enterprise Storage Server is an ideal disk solution for multiple server attachment. It provides a way to consolidate Serial disk storage into a single, powerful system that attaches to a large number of heterogeneous servers through SCSI, FC, or ESCON, and delivers many levels of advanced function.

The IBM Fibre Channel RAID Storage Server is also a new addition. It provides a robust, midrange storage solution with a native Fibre Channel host attachment. It is for servers that require both high performance and high availability in a single solution, thus providing industry standard Fibre Channel-Arbitrated Loop (FC-AL) attachment technology.

13.2.2 Disk Systems Summary

Table 132 summarizes the key features of the 2102 Fibre Channel RAID Server, 7133 Models, Enterprise Storage Server, and 2104 Expandable Storage Plus server.

Table 132. IBM Open Systems Storage External Disk Systems

External Disk System	2102-F10	7133 -T40 -D40	2105 Enterprise Storage Server	2104 Expandable Storage Plus -TL1 -DL1
Primary Usage	Performance and FC-AL connectivity	Performance	Multi-attached servers	Low COst Disk)
Drive Capacity (GB)	9.1, 18.2, 36.4	9.1, 18.2, 36.4	9.1, 18.2, 36.4	9.1, 18.2, 36.4
Entry Capacity (GB)	36.4	36.4	420	9.1
System Disk Capacity	2.2 TB	3.5 TB (per host adapter)	11.2 TB	360 GB
Attachment	FC-AL	SSA, Ultra SCSI, or FC-AL via gateways	Ultra SCSI, FC	Ultra2 SCSI
Packaging	Rack-mount	Deskside Rack-mount	Rack-mount	Tower Rack-mount

External Disk System	2102-F10	7133 -T40 -D40	2105 Enterprise Storage Server	2104 Expandable Storage Plus -TL1 -DL1
Features	<ul style="list-style-type: none"> -FC-AL attachment -multiple host attachment -RAID 0,1,3¹,5,0+1 -Distances up to 10 km -Automatic I/O path failover -Battery backed read/write cache OS support: Windows NT, Novell NetWare, AIX, HP-UX, Solaris 	<ul style="list-style-type: none"> -Serial Storage Architecture -Hot-swappable disk drives -Redundant, hot-swappable power and cooling -Up to 128 disk drives per adapter -RAID 5 using SSA RAID adapter⁴ - RAID 0, 0+1, 1 and 5 via SSA adapters -Attaches to supported models of Sun, HP, and DEC systems 	<ul style="list-style-type: none"> -Serial Storage Architecture -Dual SMP processors -Redundant power and cooling -Web-based, network-accessible storage management software -RAID 5 storage server -Attaches to supported models of UNIX, Intel-based, AS/400 and S/390 servers 	<ul style="list-style-type: none"> Low cost disk storage; Capacity from 9.1GB to 360 GB; hot-swappable drives; Redundant power and cooling option; option for second SCSI port; RAID support with SCSI RAID adapters
¹ RAID 3 capability is not supported for RS/6000 systems				

13.2.3 IBM 2102 Fibre Channel RAID Storage Server

Figure 71 shows the IBM 2102 Fibre Channel RAID Storage Server.



Figure 71. IBM 2102 Fibre Channel RAID Storage Server

The IBM 2102 is the Fibre Channel RAID Storage Server (FC RAID Storage Server). It is a high-performance, high-availability storage solution for servers running Microsoft Windows NT, Novell NetWare, AIX, HP-UX, or Sun Solaris. Its industry standard Fibre Channel host attachment provides flexibility to configure at distances of up to 10 kilometers. Dual active controllers offer a wide choice of RAID protection (0, 1, 3, 5, and 0+1) as well as redundancy. Each controller supports up to 256 MB of battery-backed, read/write cache. This is ideal for business critical applications where continuous availability and high performance are required. The IBM FC RAID Storage Server (2102-F10) supports the IBM Expandable Storage Unit (2102-D00) that provides scalability to grow to over two terabytes of data.

Characteristics

- IBM 2102 Model F10 Fibre Channel RAID Storage Server is a rack mountable dual controller unit.
- It requires 4U space.
- Can be installed in 2101-100 IBM Seascope Solutions Rack, or equivalent industry-standard 19 inch rack.
- Up to six 2102-D00 Expandable Storage Units can be attached.
- Two FC-AL ports.
- Capacities from 36.4 GB up to 2.2 TB.

- RAID protection (0, 1, 3, 5, and 0+1) with global hot sparing.
- RAID 5 arrays can span from 3+p to 19+p configurations.
- Each controller supports up to 256 MB of battery-backed, read/write cache.
- Standard HA features, such as hot pluggable power supplies and cooling units.

Limitations

- Only hosts with the same operating system can be attached to the FC RAID Storage Server when using RAID controller failover.
- A maximum of two RS/6000 system initiators and two FC RAID Storage Servers can be attached to a single FC-AL.
- RAID 3 capability is not supported for RS/6000 systems.

For additional information and supported servers, connect to the following:

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

Table 133 on page 429 summarizes the key features of the 2102 IBM Fibre Channel RAID Storage Server.

Table 133. 2102 IBM Fibre Channel RAID Storage Server

Disk Storage Capacity	(Min. - Max. GB per disk system or adapter)			
	Disk drive capacity	9.1 GB	18.2 GB	36.4 MB
System Capacity	36.4 GB up to 2.2 TB			
RS/6000 Hardware Requirements	RS/6000 Gigabit Fibre Channel Adapter for PCI Bus (# 6227) firmware level 2.22X1, or later			
Operating Systems	AIX 4.3.3 with APAR IY05369			
Supported UNIX-based Servers	SPARCserver 1000/1000E SPARCserver 2000/2000E Enterprise 2, 150, 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500, and 10000 HP 9000 Series 800, K-Class, T-Class, V-Class			
Supported Intel-based Servers	IBM Netfinity 5000 and above Pentium Pro or later, 200 MHz or faster, with 64 MB of memory or greater			
HACMP (AIX only)	HACMP 4.3.1 with APAR IY05196			
Warranty Period				
Dimensions	Width:	445 mm (17.5 in)		
	Depth:	635 mm (25.0 in)		
	Height:	175 mm (6.88 in)		
	Weight:	36 Kg (79 lb)		
Capacity of Exhaust	max. 9250 BTU/hr			
Ordering Highlights Base configuration consists of: Dual controllers with failover protection 256 (2 x 128) MB read/write cache Two FC-AL ports Six ULTRA2 SCSI LVD ports RAID 0, 1, 3, 5, and 0+1 Dual power supplies and cooling units HP,SUN,NT, Novell, and RS/6000 host support Rack mount hardware				
Feature Code	Description			
2102-F10	IBM Fibre Channel RAID Storage Server Model F10			
2102-D00	IBM Expandable Storage Unit			

The IBM Expandable Storage Unit (2102-D00) supports IBMs high-performance, reliable Ultrastar disk drives in capacities of 9.1, 18.2, and 36.4 (all 10,000 rpm) GB. The IBM Fibre Channel Switch (2109) and the IBM Fibre Channel Storage Hub (2103) provide flexibility when configuring multiple host attachments and storage servers and provides support for distances beyond 500 meters. This combination of technologies provides high performance, investment protection, and industry-leading scalability.

Included with the FC RAID Storage Server is the Fibre Channel Storage Manager (for a selected operating system) and optionally the StorWatch Fibre Channel RAID Specialist. These tools provide the ease of setup and management of the FC RAID Storage Server, and with the FC Specialist, allows multiple FC RAID Servers to be managed over a networked environment from a single workstation.

13.2.4 IBM 7133 Serial Disk System Advanced Models T40 and D40

Figure 72 shows the IBM 7133 Serial Disk System Advanced Models T40 (deskside tower) and D40 (rack-mountable drawer).



Figure 72. IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer

These IBM 7133 Serial Disk System models are leading-edge serial storage systems that use industry-standard Serial Storage Architecture (SSA) to deliver outstanding performance and capacity/footprint, improved connectivity, availability, and enhanced attachability.

The IBM 7133 Model D40 is a four Electronics Industries Association (EIA) units high drawer designed to be mounted in a standard 19-inch rack. The IBM 7133 Model T40 is a stand-alone, deskside tower unit.

Both models have available slots for up to sixteen Advanced Disk Drive Modules. These sixteen disks are clustered together in groups of four, such as the drives in the previously announced IBM 7133 models 010, 020, 500, and 600. Also consistent with the 7133 models 020 and 600, the models D40 and T40 provide circuitry that automatically connects these groups of disk drives. D40 and T40 models support SSA160 or SSA80 speeds for up to 160

MB/s data transfer. Earlier models of 7133s and other SSA systems support SSA80 speeds.

Each model can be populated with Advanced Disk Drive Modules with capacities of 9.1 GB, 18.2, or 36.4 GB. Any combination of drive sizes can be intermixed to populate a 7133 with 4 to 16 disks.

The 7133 has multiple, selectable cable features from 1 to 25 meters long for attaching the Disk System to other SSA devices. The Advanced SSA Cables used with models D40 and T40 provide enhanced electrical characteristics over previous cables and can be recognized by their blue color. A minimum of two cables are required to properly install the 7133 in an SSA loop. However, more cables may be needed to achieve the desired performance and availability characteristics.

These models introduce a new fiber optic extender feature, the Advanced SSA Optical Extender, which can connect D40 or T40 models or SSA 160 adapters at distances up to 10 km. When connected to other SSA devices (other 7133 models or SSA adapters), the Advanced SSA Optical Extender supports distances up to 2.4 km.

As an option, the 7133 supports remote power control from up to two host systems.

A variety of performance and data protection methods, such as Fast Write Cache, RAID 0 striping, RAID 1 mirroring, and RAID 5 parity protection, are supported by the 7133 using specific adapter or host software facilities.

Characteristics

- No single point of data-path failure.
- Up to sixteen 9.1, 18.2, or 36.4 GB hot-swappable SSA disk drive modules.
- Provides high capacity--up to 582 GB per tower or drawer and 3.5 TB per host adapter.
- Concurrent maintenance of redundant power and cooling.
- Connects up to eight servers.
- Up to 10 km between disk systems with the Advanced SSA Optical Extender.
- Compliant with the SSA interface specifications (ANSI X3T10.1).
- Enables disk sharing through simultaneous attachment of multiple UNIX and Windows NT hosts to the IBM 7133 Serial Disk System using the

2108-S20 IBM Storage Area Network Data Gateway for Serial Disk or the 7139 Vicom FC SLIC Router.

- Attaches to selected Intel-based servers using the IBM SerialRAID/X Adapter.
- Models D40 and T40 may be intermixed in an SSA loop with other 7133 models and the 7131-405.
- New D40 power supply option (-48 V DC) is available for use in the telecommunications industry.

Table 134 summarizes the key features of the 7133 Serial Disk System Advanced Models D40 and T40.

Table 134. IBM 7133 Serial Disk System Advanced Models at a Glance

Disk Storage Capacity	(Min-max GB per disk system or adapter)			
	Disk drive capacity	9.1 GB	18.2 GB	36.4 GB
	Per disk system	36.4 - 145.6	72.8-291.2	145.68-582.4
System Capacity	Up to eight disk systems per adapter Up to 128 drives per adapter			
Currently Available RS/6000 Adapters	# 6215 SSA Multi-initiator RAID Enhanced Loop Adapter # 6230 Advanced SerialRAID Plus Adapter			
Operating Systems	AIX Version 4.3, or later The specific AIX release level requirements are RS/6000 host system and SSA adapter dependent.			
Other Supported UNIX-based Servers	Sun SPARCstation 20; SPARCserver 1000 and 1000E SPARCcenter 2000 and 2000E Ultra Enterprise 2, 3000, 4000, 5000, and 6000 SPARCstation 10 and 20 SPARCserver 1000 and 1000E SPARCcenter 2000 and 2000E Ultra Enterprise 2, 30, 60, 150, 450, 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500, 10000 HP 9000 Enterprise Servers D, K, T, G, E, H, I, V class Compaq (Digital) AlphaServer 300, 800, 1000, 1000A, 2000, 2100, 2100A, 4000, 4100, 8200, 8400 Series and GS140			
Supported Intel-based Servers	IBM PC Server models 325 and 330 IBM Netfinity Servers Model 5000, 5500, 5500-M10, 5500-M20, 5500-Xeon, 7000, 7000-M10 5600, 8500R Compaq ProLiant Servers 5500, 6500-Xeon, 7000-Xeon HP NetServer Models LH2 Dell PowerEdge 4300, 6350			

Warranty Period	One-year limited warranty	
Dimensions	Model D40	Rack mountable drawer, 4EIA units 6.7"H x 17.5"W x 29"D (171 mm x 444 mm x 737 mm)
	Model T40	Deskside Tower 24.0"H x 8.3"W x 33.9"D (610 mm x 210 mm x 861 mm)
Ordering Highlights Base unit does not include disk drives. Add any of the supported disk drive modules for a total of four to sixteen disk drives. Two power supplies are required.		
Feature Code	Description	
8031	Raven-black cover for model D40	
8032	Pearl-white cover for model D40	
8509	10K/9.1 GB disk drive	
8518	10K/18.2 GB disk drive	
8436	36.4 GB disk drive	

For additional information, contact the following Web site:

<http://www.storage.ibm.com/hardsoft/products/7133/>

In High Availability configurations, the Model D40 and T40 are supported by:

- HACMP V4.3.0
- HACMP V4.2.2

13.2.5 Enterprise Storage Servers (2105)

This section describes the Enterprise Storage Server (ESS) family and includes following models:

- 2105-E10 IBM Enterprise Storage Server Model E10
- 2105-E20 IBM Enterprise Storage Server Model E20
- 2105-F10 IBM Enterprise Storage Server Model F10
- 2105-F20 IBM Enterprise Storage Server Model F20

The model summary matrix is shown in Table 135. The E20 and F20 models support an expansion frame to provide greater capacity over the E10 or E20 models.

Table 135. Model Summary Matrix

Model Number	Power	App. End-User Capacity	Model Conversion
2105-E10	Single Phase	420 GB - 1.7 TB	2105-F10
2105-E20	Three Phase	420 GB - 11.1 TB	2105-F20
2105-F10	Single Phase	420 GB - 1.7 TB	
2105-F20	Three Phase	420 GB - 11.1 TB	

e-business is generating exponential growth in immediately accessible, always available, and highly functional storage capacity. Dependable, appliance-like intelligent storage, served by its own network, is required to provide information as a utility.

The Enterprise Storage Server combines and integrates multiple storage technologies into a single disk system. Added-value functions, such as remote mirroring already available in IBMs popular S/390 RAMAC disk system family, are extended in the ESS to support platforms, such as NetWare, Windows NT, Windows 2000, and UNIX.

Enterprise Storage Server Highlights

- Comprehensive Storage Consolidation Platform
 - Scalable to over 11 TB usable RAID-protected capacity
 - ESCON, Fibre Channel, UltraSCSI, FICON (preview) connections
 - Up to 32 direct host connections; hundreds possible in a SAN
 - Choice of disk capacities, most 10K RPM
- Breakthrough Performance
 - Large cache
 - Internal bus bandwidth near 800 MB/s
 - Up to 64 internal disk paths / all can be transferring data concurrently
 - Data striped for balanced disk utilization and parallel access
 - Specialized S/390 performance accelerators
- Round-the-Clock Access to Data
 - No hardware single points of failure

- Concurrent maintenance, repairs, microcode changes, upgrades supported
- Hot-pluggable / hot-swappable components
- Storage Area Network (SAN) Support
 - Fibre Channel attachment, including switch, hub, and gateway support
- Added-Value Capabilities
 - Near-instantaneous volume copy (FlashCopy)
 - Remote Copy (remote mirroring) options (PPRC and S/390 XRC)

The IBM Enterprise Storage Server is a superior Storage Area Network (SAN) storage solution that:

- Supports critical requirements for strategic business initiatives including:
 - e-business
 - Enterprise resource planning
 - Business intelligence
 - Server consolidation
 - All other mission-critical applications
- Deploys the IBM Seascope principles:
 - Powerful storage server
 - Snap-in building blocks
 - Universal data access

The physical configuration of the Enterprise Storage Server is composed of the following components:

- Storage Server - Composed of two RISC servers
- Cluster - Provides the facilities to control and manage data transfer and advanced functions. Should one cluster fail, the remaining cluster can take over the functions of the failing cluster. A cluster is composed of the following sub-components:
 - Host Adapters - Each cluster has one or more host adapters (HAs). Each host adapter provides one or more host I/O interfaces. A host adapter can communicate with either cluster complex.
 - Device Adapters - Each cluster has one or more device adapters (DAs). Each device adapter provides one or more storage device interfaces. Disk drives are attached to a pair of device adapters, one in

each cluster, so that the drives are accessible from either cluster. At any given time, a given disk drive is managed by only one device adapter.

- **Cluster Complex** - The cluster complex provides the management functions for the Enterprise Storage Server. It consists of cluster processors, cluster memory, cache, nonvolatile storage (NVS), and related logic.
 - **Cluster Processor** - The cluster complex contains four cluster processors (CP) configured as symmetrical multiprocessors (SMP). The cluster processors execute the licensed internal code that controls operation of the cluster.
 - **Cluster Memory, Cache** - Is used to store instructions and data for the cluster processors. The cache memory is used to store cached data from the disk drives. The cache memory is accessible by the local cluster complex, by device adapters in the local cluster, and by host adapters in either cluster.
 - **Nonvolatile Storage (NVS)** - Is used to store a nonvolatile copy of active written data. The NVS is accessible to either cluster-processor complex and to host adapters in either cluster. Data may also be transferred between the NVS and cache.
- **Disk Drives** - Provide the primary nonvolatile storage medium for any host data stored within the Enterprise Storage Server. Storage devices are grouped into ranks and are managed by the clusters.

Complete and current information on all supported servers, server attachment adapters, and operating systems are available on:

<http://www.storage.ibm.com/hardsoft/products/ess/supserver>

Table 136 lists the RS/6000 systems, which could be attached to the ESS.

Table 136. List of RS/6000 Systems

Models	Operating Systems	SCSI Host Adapters	SCSI Cables
7025-F50	AIX 4.2.1, AIX 4.3	# 6207	# 9701, # 9702
7025-F80	AIX 4.3.3	# 6207	# 9701, # 9702
7026-H70	AIX 4.3	# 6207	# 9701, # 9702
7026-H80	AIX 4.3.3	# 6207	# 9701, # 9702
7026-M80	AIX 4.3.3	# 6207	# 9701, # 9702

Models	Operating Systems	SCSI Host Adapters	SCSI Cables
7017-S7A	AIX 4.3	# 6207	# 9701, # 9702
7017-S80	AIX 4.3.3	# 6207	# 9701, # 9702
9076	AIX 4.2.1, PSSP 2.4 AIX 4.3, PSSP 3.1	# 6207 # 6209	# 9701, # 9702 # 9703, # 9704

Enterprise Storage Server attachment to storage area networks or host systems using an industry-standard Fibre Channel Arbitrated Loop (FC-AL) interface is provided through the IBM 2108 Storage Area Network (SAN) Data Gateway Model G07.

Table 137 lists Enterprise Storage Server characteristics at a glance.

Table 137. ESS Characteristics at a Glance

Characteristics	ESS 2105-E10/2105-F10	ESS 2105-E20/2105-F20
Disk storage capacity	420 GB to 1.68 TB	420 GB to 3.36 TB
Max. storage system capacity	1.68 TB	11.2 TB
Cache size	6 GB	6 GB
Host server attachments	Up to 32 ESCON or SCSI ports; Fibre Channel connection via IBM SAN Data Gateway	Up to 32 ESCON or SCSI ports; Fibre Channel connection via IBM SAN Data Gateway
Supported Systems ¹ : S/390, AS/400 (9406 models) Data General; Compaq, HP (9000 and 8000), Intel-based PC servers; Novell NetWare; Sun, RS/6000, RS/6000 SP		
Dimensions	75.25" H x 54.50" W x 35.75" D (1913 mm x 1383 mm x 909 mm)	75.25" H x 54.50" W x 35.75" D (1913 mm x 1383 mm x 909 mm)
Caloric Value	20.000 BTU/hr	34.000 BTU/hr
Electrical Power	6.0 kVA	10.1 kVA

HACMP Hardware Support for ESS Models F10 and F20

HACMP for AIX, Versions 4.3.1 and 4.4 now support the following previously announced hardware. Support for ESS Models F10 and F20 with SCSI attachment or Fibre attachment using the IBM SAN Data Gateway Model G07(2108-G07) is available (Table 138 on page 438). Direct Fibre attachment is not yet supported by HACMP to any model of the ESS.

For additional information see the following web page:

<http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm#6>

Table 138. HACMP hardware support for ESS models

HACMP/ESS Hardware and Software Support	HACMP 4.2.2 AIX 4.2.1	HACMP 4.2.2 AIX 4.3.3	HACMP 4.3.1 AIX 4.3.3	HACMP 4.4 AIX 4.3.3
IBM Enterprise Storage Server 2105-F10, 2105-F20	Not supported	Not supported	HACMP APAR IY11110 IY11560 IY11564	HACMP APAR IY11480 IY11563 IY11565
IBM SAN Data Gateway Model G07 2108-G07 Fibre Attachment	Not supported	Not supported	Note 2, HACMP APAR IY07313	Note 2, HACMP APAR IY10564 (Note 3)
IBM Enterprise Storage Server 2105-E10, 2105-E20	HACMP APAR IY04403 (Note 1)	HACMP APAR IY04403 (Note 1)	HACMP APAR IY03438 (Note 1) IY11560 IY11564	HACMP APAR IY11563 IY11565
IBM Data Path Optimizer for Unix Version 1.1 5648-B58 (SDD 1.1.4) SCSI Attachment	Not supported	Not Supported	Note 4, HACMP APAR IY07392	Note 4
<p>1. Required for the CRM or ESCRM feature codes only.</p> <p>2. This now includes the features 2214, 2313, and 2319. When using two fibre channel adapters to a single SAN Data Gateway in an HACMP environment, the SAN Data Gateway must be run in "split mode". Each fibre channel adapter must be connected to a separate port on the SAN Data Gateway and two separate ports on the SAN Data Gateway must be used to attach to two separate SCSI ports on the ESS. The SAN Data Gateway will not be supported in a switched environment when HACMP support for switched fibre channel becomes available.</p> <p>3. Required when attaching to the ESS model E10.</p> <p>4. DPO is only supported in concurrent mode available in the CRM and ESCRM features of HACMP. For ESS installations, use Subsystem Device Driver Version 1.1.4 in place of DPO. DPO Version 1.1 and SDD Version 1.1.4 are not supported with the SAN Data Gateway</p>				

13.2.6 Expandable Storage Plus (# 2104)

The IBM Expandable Storage Plus disk enclosure provides low cost disk storage for RS/6000 servers. In addition, Expandable Storage Plus offers space for up to ten hot swap Ultra 2 (LVD) disk drives, including IBM's high performance 9 GB, 18 GB, and 36 GB Ultrastar drives. Figure 73 shows the IBM Expandable Storage Plus Model DL1 and Model TL1.



Figure 73. IBM Expandable Storage Plus SCSI

It is designed to meet the needs of RS/6000 customers that have external disk capacity requirements and desire to use a Ultra 2 (LVD) interface. When combined with the IBM PCI 3-channel Ultra SCSI RAID Adapter, high performance and high availability are available. IBM's Expandable Storage Plus provides leading edge technology at a very affordable price.

IBM supports HACMP for the IBM Expandable Storage Plus when used with RS/6000 servers and non-RAID adapters.

13.2.6.1 Supported Adapters

The IBM Expandable Storage Plus attaches to RS/6000 servers with the following adapters:

- IBM SCSI-2 Fast/Wide Adapter/A (# 2415)
- IBM PCI RAID Adapter (# 2493)
- PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)
- PCI Dual Channel Ultra2 SCSI Adapter (# 6205)
- PCI Single-Ended Ultra SCSI Adapter (# 6206)

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

The appropriate cable to connect the adapter and Expandable Storage Plus is also required.

13.2.6.2 At a Glance

The IBM Expandable Storage Plus provides:

- Disk Enclosure for up to 10 high-performance hot swap IBM Ultra 2 (LVD) disk drives, including both 7,200 and 10,020 rpm options
- Capacity ranging from 9 GB to 360 GB, utilizing 9, 18, or 36 GB disk drives
- Redundant power and cooling
- Advanced monitoring using SCSI Enclosure Services
- Rack and Tower models
- Support across a wide variety of RS/6000 servers
- IBM SystemXtra support services and financing

Table 139 lists the current available RS/6000 systems.

Table 139. Available Products Matrix for the IBM Expandable Storage Plus

RS/6000 Server	Adapter	AIX Level
7043-150	FC 2494, 6205, 6206	4.2.1, with PTFs, 4.3.3, or later
7044-170	FC 2494, 6205, 6206	4.3.3, or later
7044-270	FC 2494, 6205, 6206	4.3.3, or later
7025-F50	FC 2493, 2494, 6205, 6206, 6208	4.2.1, with PTFs, 4.3.3, or later
7025-F80	FC 2494, 6204, 6205	4.3.3 or later
7026-H70	FC 2493, 2494, 6205, 6206, 6208	4.2.1, with PTFs, 4.3.3, or later
7026-H80	FC 2494, 6204, 6205	4.3.3, or later
POWER3 SMP High Node	FC 6205	4.3.3, or later
SP I/O Expansion I/O	FC 6205	4.3.3, or later

Refer to the following Web site to obtain the latest information:

<http://www.storage.ibm.com/hardsoft/products/expplus/supserver.htm>

13.2.7 SAN/Fibre Channel Interconnection

This section is divided into two subsections:

- SAN Data Gateway and Router Family
- IBM SAN Fibre Channel Switch and Storage Hub

13.2.7.1 SAN Data Gateway Family

The IBM 2108 Storage Area Network (SAN) Data Gateway is a family of connectivity products that enable fibre channel connections to Ultra-SCSI or SCSI-based storage or host systems. The family includes:

- 2108-G07 IBM Storage Area Network Data Gateway Model G07
- 2108-R03 IBM SAN Data Gateway Router Model R03
- 2108-S20 IBM Storage Area Network Data Gateway for Serial Systems Model S20

13.2.7.2 2108-G07 SAN Data Gateway

Figure 74 shows the SAN Data Gateway (2108-G07).




Figure 74. IBM SAN Data Gateway

The IBM Storage Area Network (SAN) Data Gateway enables you to attach SCSI and Ultra SCSI storage servers and storage devices to Fibre Channel environments. It attaches new or existing SCSI storage products to storage area networks or directly to host systems using an industry standard Fibre Channel Arbitrated Loop (FC-AL) interface. The SAN Data Gateway is a protocol convertor that is:

- Easy to manage
- Configures with up to three FC-AL ports
- Provides four Ultra SCSI Differential ports for device attachment

The SAN Data Gateway provides access between its fibre channel ports and its SCSI ports. Channel zoning provides access control between ports. While channel zoning provides control of paths between host adapters and SCSI storage ports, it does not limit access to specific devices (LUNs) within the storage system. Virtual Private SAN (VP SAN) provides LUN masking to limit

access between host adapters and LUNs attached to SAN Data Gateway SCSI ports. The IBM StorWatch SAN Gateway Specialist, an easy to use graphical user interface, provides the tools to define SAN Data Gateway channel zoning, the VP SAN LUN-masking, and controls which host systems have access to specific storage devices.

Multiple Configuration Options

The SAN Data Gateway utilizes Fibre Channel and Ultra SCSI channel bandwidth for high-performance attachment of the following devices:

- IBM Enterprise Storage Server
- IBM Magstar 3590 Tape Subsystem in stand-alone, Magstar 3494 Tape Library, and Magstar 3590 Silo Compatible Tape Subsystem environments
- IBM Magstar MP 3570 Tape Subsystem or Magstar MP 3575 Tape Library Dataserver
- IBM 7337 DLT Library

The SAN Data Gateway is available as a rack-mounted unit or as a stand-alone tabletop unit. The full-function gateway model provides VP SAN, two short-wave Fibre Channel ports, and four Ultra SCSI Differential ports to attach disk or tape storage devices. The low-cost Router model provides one short-wave Fibre Channel port and two Ultra SCSI Differential or Single-End SCSI ports for tape storage.

The SAN Data Gateway is a compact, self-contained unit. There are four standard Ultra SCSI Differential ports, each with a single 68-pin HD connector for attachment of storage servers and devices. The SAN Data Gateway can be configured with up to three FC-AL ports.

For a single FC-AL port configuration, specify feature number 2204. For a two port configuration, specify feature number 2205. For a three port configuration, specify feature number 2205 and feature number 2303 (Quantity = 1).

A SAN Data Gateway initially configured with one FC-AL port (# 2204) can be field upgraded to a two or three port configuration by ordering one or two additional FC-AL ports (# 2303). Gateways initially configured with two FC-AL ports (# 2205) can be field upgraded to a three port configuration by ordering one additional FC-AL port (# 2303). Initial plant orders for three port configurations should order the two port feature (# 2205) and one additional FC-AL port (# 2303). Each FC-AL port provides a duplex shortwave optical interface, and supports attachment of shortwave fiber optic cables with an

industry standard duplex SC connector to specified host adapters or to other storage area network components, such as Fibre Channel hubs.

13.2.7.3 2108-R03 IBM SAN Data Gateway Router

The SAN Router is a compact, self-contained unit. There is one standard Fibre Channel shortwave optical port for attachment to Fibre Channel-capable host systems or other SAN Fibre Channel connectivity products and two standard Ultra SCSI ports. By selecting # 2820, or ordering part number 2108R3S, the two Ultra SCSI ports are single-end SCSI. If you select # 2830, or order part number 2108R3D, the two Ultra SCSI ports are high-voltage differential.

The IBM Storage Area Network Data Gateway Router (2108-R03) mounts in the IBM Seascope Solution Rack (2101-100), the Netfinity Rack (9306900), or an equivalent, industry standard 19-in. rack, or, optionally, can be configured as a stand-alone tabletop unit.

In a rack-mounted configuration, 2U space is required, and one outlet on the rack Power Distribution Unit (PDU) is required. The power cord for attachment to the rack PDU is provided. For stand-alone tabletop configurations, you must specify feature number 9205 and specify the country-unique power cord feature number (98xx) or order the appropriate country-specific power cord part number.

The IBM SAN Data Gateway Router contains a self-regulating power supply that connects to an electrical source of 100v to 240v at 50 or 60Hz. Included with the SAN Data Gateway is a single power cord for attachment to the power distribution unit (240v) with the IBM Seascope Solutions Rack (2101-100) or equivalent.

When feature number 9205 is selected, a single country-specific power cord will be supplied based on the 98xx feature number specified (U.S./Chicago use # 9986).

13.2.7.4 7139-111 Vicom Systems SLIC Router FC-SL

The 7139 Vicom Systems SLIC Router FC-SL is one of Vicoms connectivity products that enables Fibre Channel based servers to connect to IBMs SSA-based disk subsystems.

The IBM 7139 Model 111 has a single fibre channel port for connection to an SSA loop consisting of IBM 7131, 7133, or 3527 subsystems, easy-to-use tools for initial setup and ongoing management, advanced functions, such as RAID-1 data protection and instant copy for improved SSA loop availability,

and ability to support large amounts of SSA-based disk storage using multiple SLIC routers.

The Vicom Systems SLIC Router FC-SL provides the following:

- Connectivity for fibre channel based servers to IBMs Serial Storage Architecture (SSA)-based disk subsystems.
- All the advantages of SSA architecture, including high performance, high reliability, and scalability.
- Vicoms Serial Loop IntraConnect (SLIC) technology, which is used worldwide.
- Easy-to-use, graphical setup and ongoing management tools.
- Advanced functions, such as RAID-1 data replication, for added data protection and Instant Copy for SSA loop availability.
- Support for large amounts of SSA-based disk storage when users take advantage of multiple SLIC routers.

The Router attaches with a single short-wave or long-wave GBIC and a Fibre Channel cable. Optionally, one additional GBIC may be ordered to extend the Fibre Channel loop to other equipment.

For more information, see the Vicom Web page at:

<http://www.vicom.com>

13.2.7.5 2108-S20 IBM SAN Data Gateway for Serial Systems

Figure 75 shows IBM Storage Area Network Data Gateway for Serial Systems Model (2108-S20).




Figure 75. IBM SAN Data Gateway Model S20

New or existing UNIX and Windows NT host systems can now take advantage of the many benefits and features of the 7133 Serial Disk systems and serial storage technology (where native attachment is not available).

The SAN Data Gateway Model S20:

- Provides multi-host, highly available data-sharing solutions - More than two terabytes of serial disk storage can be attached on a single Gateway S20 loop, and up to eight Gateway Model S20s can be configured on a loop.
- Offers increased data availability, performance, and protection - The SAN Data Gateway Model S20 lets you replicate data across serial disk systems and simultaneously mirror two or three copies of data, independent of the host system.
- Includes the StorWatch SAN Data Gateway S20 Specialist, a NT based management tool that provides configuration and service functions. The Specialist enables you to define RAID 1 groups, disk concatenation, and supports splitting off a copy of a mirrored group for backup or processing on another system.

The SAN Data Gateway Model S20 is a compact, self-contained unit. It provides one wide Ultra SCSI Differential port and two SSA ports. The UltraSCSI Differential port has two 68-pin HD female connectors. An external terminator is provided on one of the SCSI connectors. Both fast/wide SCSI-2 Differential and fast/wide UltraSCSI Differential are supported. The Gateway will negotiate to the connected host adapter interface, complying with ANSI SCSI-2 specification X3.131-1994.

The two SSA ports each have a 9-pin, miniD female connector for attaching to a single SSA loop containing 7133 Serial Disk Systems, 7131 Model 405 Multi-Storage Towers, and 3527 SSA Storage Systems.

The SAN Data Gateway Model S20 can be used as a stand-alone table top unit or be mounted in a Seascope Solution Rack (2101-100) or equivalent (7014-S00, 7015-R00, 7202-900). For rack mounting, the optional feature 2000 must be ordered. This feature provides a 2U rack-mountable enclosure that can hold up to four Gateway model S20s. Included with the rack enclosure are two bifurcated power cords that connect the rack power distribution unit (PDU; 240V) within the 2101-100 rack or equivalent. Two outlets are required on the PDU. Also supplied with each Gateway S20 is a single country-specified wall outlet power cord based on the 98xx feature number selected (US/Chicago use # 9986). The Gateway contains a self-regulating power supply that connects to an electrical source, 100V to 240V at 50 to 60HZ.

The SAN Data Gateway requires a SCSI cable for attaching to a host server and two SSA cables for attaching to the Serial Disk Systems. These cables are customer supplied and typically ordered with the attaching host server

and Serial Disk systems. The total SCSI bus length cannot exceed the differential SCSI limit of 25 meters. The host SCSI bus must be connected only between the host bus adapter and the SAN Data Gateway S20. Table 140 shows the SAN Data Gateway for Serial Disk at a glance.

Table 140. IBM SAN Data Gateway for Serial Disk at a Glance

Characteristics	2108 Model S20
Storage systems	Attaches to IBM 7133 Serial Disk System Advance Models D40 and T40 (SSA160) and all SSA80 models
Supported servers and operating systems	Intel-based Netfinity and Compaq servers with Windows NT 4.0 Server or Enterprise Server Edition Sun SPARCserver and Enterprise Server Edition with Solaris 2.5.1, 2.6, or 2.7 HP 9000 Series servers with HP/UX 10.10, 10.20, or 11.0 Compaq (Digital) AlphaServers with Tru64 UNIX 4.0
Stand-alone unit dimensions	76.0 mm H x 107.0 mm W x 279.4 mm D (3.00" x 4.26" x 11.00")
Drawer option dimension	88.9 mm H x 483.0 mm W x 507.0 mm D (3.50" x 19.00" x 19.95")
Electrical power requirements	60 Watts

13.2.8 IBM SAN Fibre Channel Switch and Storage Hub

The IBM 2109 SAN Fibre Channel Switch family includes two members:

- IBM 2109 Model S08
- IBM 2109 Model S16

The IBM 2109 SAN Fibre Channel Switch is used to interconnect multiple host servers with storage servers and devices, thus creating a storage area network (SAN).

The IBM SAN Fibre Channel Switch offers:

- Industry-standard Fibre Channel host, storage, and storage area network attachment
- Scalability from small- to very large-size SAN environments, providing scalable bandwidth
- Robust performance that handles large-block transfers very fast and reliably, with high data integrity
- Automatic configuration of simple Fabric, zoning control for more complex Fabrics

- Modularity for flexible system configurations, including diagnostics to isolate problems quickly
- An embedded Web browser interface for configuration, management, and service
- Support for Public Fibre Channel Arbitrated Loops

The IBM 2109 Model S08 has an 8-port Fibre Channel switch and four shortwave Gigabit Interface Convertors (GBICs).

The IBM 2109 Model S16 has a 16-port Fibre Channel switch and four shortwave Gigabit Interface Convertors (GBICs).

The IBM 2103 is a Fibre Channel Storage Hub that can be used to ease the configuration of Fibre Channel Arbitrated Loops and to add greater distances to the FC-AL topology.

The IBM 2103 Model H07 Fibre Channel Storage Hub can be used to ease the configuration of Fibre Channel Arbitrated Loops and to add greater distances to the FC-AL topology. The FC Storage Hub provides four short wave optical Gigabit Interface Converter (GBIC) ports and the option to add up to three additional long wave or short wave optical GBIC ports. The short wave GBIC ports allow fiber cable connection of up to 500 meters to the host based FC-AL initiator adapter to the FC-AL port on the FC RAID Storage Server or the FC port on the SAN Data Gateway or Router (2108-G07/R03). Either short wave or long wave GBIC ports can be used to connect two FC Storage Hubs, therefore, extending the distance up to an additional 500 meters with the short wave GBIC or up to an additional 10 kilometers with the long wave GBIC. The IBM Fibre Channel Storage Hub comes in a rack mounted tray that requires 1U of rack space. If two FC Storage Hubs are ordered, they will both be assembled in the same 1U rack mounted tray. An option (# 9205) is available for a non-rack mounted, tabletop usage with an external country-unique power cord.

13.2.9 IBM 7204 External Disk Drives

Figure 76 shows an IBM external disk drive.



Figure 76. IBM 7205 External Disk Drive

To add external storage capacity to your RS/6000, simply attach a self-powered drive. There are two models, and Table 141 summarizes the key features of the 7204 External Disk Drives.

Table 141. IBM 7204 External Disk Drives at a Glance

Model	GB	Interface	H/W (#) Requirements
409	9.1	SCSI-2 FW SE	6205, 6208
419	18.2	SCSI-2 FW SE	6205, 6208
Dimensions			2.36"H x 9.84"W x 10.83"D (60 mm x 250 mm x 275 mm)
Operating Systems	AIX 4.3		

Note

To connect a # 6205 to the following devices:

- 7204-409 External 9.1 GB disk drive
- 7204-419 External 18.2 GB disk drive
- 7206-220 External 4 mm tape drive
- 7332-220 External 4 mm tape drive

You need to order a no-charge VHDCI Cable/Interposer (# 9799) and one of the appropriate cables:

- 9137 1.0 m Ultra2 SCSI cable
- 9149 1.0 m SCSI-2 F/W cable

13.2.10 IBM 7203 Portable Disk Drive

The IBM 7203 External Portable Disk Drive is a self-powered external disk drive that provides additional disk storage for the RS/6000 family of systems. It includes a removable disk drive module that provides portability and enhanced data security.

It attaches to a RS/6000 SCSI I/O (input/output) controller with a narrow, single-ended interface. Several 7203 Disk Drives can attach simultaneously to the SCSI I/O controller.

The 7203 includes a removable disk module that provides portability and enhanced data security. The data modules can be easily removed and transported to move data files from host to host system to enable:

- Updating user programs
- Transporting large data files
- Adding new software applications

The module can contain disk drives of various capacities. Additional modules can be purchased for increased storage.

The 7203 model 001 may be configured with the following option:

- 9.1 GB Select Option (# 2315) and 9.1 GB Portable Disk Module (# 2350)

This replaces the standard 355 MB Portable Disk Drive Module with the 9.1 GB Portable Disk Module on the initial plant order of the IBM External Portable Disk Drive (7203-001).

13.3 IBM Tape Drive Products

In this section, the following tape drive products are discussed:

- 7205 Digital Linear Tape Drive
- 7206 4 mm DDS-2, DDS-3 and DDS-4 Tape Drives
- 7207 1/4" Tape Drives
- 7208 8 mm Tape Drive

First, the key features of these tape drives are summarized in Table 142; more detailed descriptions follow in subsequent sections.

Table 142 summarizes the external tape drives that are available from IBM.

Table 142. IBM Tape Drive Products

Machine Model	7206 -005 -110 -220	7207 -122	7208 -341	7205 -311
Description	4 mm DDS-2, DDS-3, DDS-4 Drive	1/4 Inch Drive	8 mm Drive	Digital Linear Tape
Number of Drives	1	1	1	1
Max. Number of Cartridges	1	1	1	1
Native Cartridge Capacity	4 GB 12 GB 20 GB	4 GB	20 GB	35 GB
Max. Native Capacity per Drive/Library	4 GB 12 GB 20 GB	4 GB	20 GB	35 GB
Max. Capacity Compressed	8 GB 24 GB 40 GB	8 GB	40 GB	70 GB
Drive Data Rate per Sec	400 KB 1.1 MB 3 MB	380 KB	3 MB	5 MB
Max. Compressed Data Rate	800 KB 2.2 MB 6 MB	760 KB	6 MB	10 MB

Machine Model	7206 -005 -110 -220	7207 -122	7208 -341	7205 -311
Interface	SCSI-2 F/W SCSI-2 F/W Ultra2 SCSI, SCSI-2 F/W	SCSI-2F/WDE, F/WSE SCSI	SCSI F/W	SCSI-2
Supported Platforms	IBM HP SUN DEC	IBM	IBM	IBM

13.3.1 IBM 7205 Digital Linear Tape Drive

Figure 77 shows the IBM 7205 Digital Linear Tape Drive.



Figure 77. IBM 7205 Digital Linear Tape Drive

The 7205 employs DLT 7000 drive technology and is designed for users who have large amounts of data to back up or limited time for system backup. It is read/write compatible with previous DLT tape technologies.

- Up to 35 GB/70 GB data storage capacity
- Up to 5 MB/s / 10 MB/s data transfer rate

Table 143 summarizes the key features of the 7205 Model 311 Digital Linear tape drive.

Table 143. IBM 7205-311 Digital Linear Tape Drive at a Glance

Characteristics	7205 Model 311
Capacity	35 GB
Capacity compressed	70 GB
Media	1/2-inch DLT tape-IV

Data transfer rate	5 MB/s
Data transfer rate compressed	10 MB/s
SCSI interface	SCSI-II differential
Operating Systems	AIX Version 4.2.1, or later
Media	59H3040
Warranty Period	One-year limited warranty
Dimensions	4.8"H x 11.0"W x 11.5"D (122.3 mm x 280 mm x 290 mm)

13.3.2 IBM 7206 External 4 mm DAT Tape Drives

Figure 78 shows an IBM 7206 external 4 mm DAT tape drive.



Figure 78. IBM 7206 External 4 mm DAT Tape Drive

These DAT cartridge tape drives feature a data grade media recognition system and host-independent diagnostics. An SCSI cable, 4 mm data cartridge, automatic diagnostic cartridge, and cleaning cartridge are included.

7206-005 DDS-2

- Up to 4 GB / 8 GB data storage capacity
- Up to 400 KB/s / 800 KB/s data transfer rate

7206-110 DDS-3

- Up to 12 GB / 24 GB data storage capacity
- Up to 1.1 MB/s / 2.2 MB/s data transfer rate

7206-220 DDS-4

- Up to 20 GB / 40 GB data storage capacity
- Up to 3 MB/s / 6 MB/s data transfer rate

Table 144 summarizes the key features of the 7206 Model 110 DDS-2 DAT tape drive.

Table 144. IBM 7206-110 DDS-3 DAT Tape Drive at a Glance

Characteristics	7206 Model 110
Tape drive type	4 mm
Capacity per cartridge	12 GB
Capacity per cartridge compressed	24 GB
Data transfer rate	1 MB/s

Data transfer rate compressed	Up to 2 MB/s
Interface	SCSI-2
Media	4 mm 125-meter media (P/N 59H3465)
Supported systems	Selected ibm RS/6000 models
Operating Systems	AIX Version 4.2, 4.3 or later
Media	59H3465
Warranty Period	One-year limited warranty
Dimensions	2.2" H x 9.8" W x 10.8" D (55 mm x 250 mm x 275 mm)

Table 145 summarizes the key features of the 7206 Model 220 DDS-4 DAT tape drive.

Table 145. IBM 7206 Model DDS-4 at a Glance

Characteristics	7206 Model 220
Capacity per cartridge	20 GB
Capacity per cartridge compressed	40 GB
Data transfer rate	3 MB/s
Data transfer rate compressed	6 MB/s
RS/6000 Hardware Requirements	SCSI-2 Fast/Wide PCI-Bus Adapter (# 2408) SCSI-2 Fast/Wide Adapter/A (# 2415) PCI Dual Channel Ultra 2 SCSI Adapter (# 6205) PCI Single-Ended Ultra SCSI Adapter (# 6206) PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)
Operating Systems	AIX Version 4.2, 4.3, or later
Media	59H4458 4 mm DDS-4 Data Cartridge 59H4457 4 mm DDS-4 Test Cartridge 21F8763 4 mm DDS-4 Cleaning Cartridge
Warranty Period	One-year limited warranty

Dimensions	2.2"H x 9.8" W x 10.8" D (55 mm x 250 mm x 275 mm)
-------------------	---

Note

To connect a #6205 to the following devices:

- 7204-409 External 9.1 GB disk drive
- 7204-419 External 18.2 GB disk drive
- 7206-220 External 4 mm tape drive
- 7332-220 External 4 mm tape drive

You need to order no-charge VHDCI Cable/Interposer (# 9799) and one of the appropriate cables:

- 9137 1.0 m Ultra2 SCSI cable
- 9149 1.0 m SCSI-2 F/W cable

13.3.3 IBM 7207 1/4" Cartridge Tape Drives

Figure 79 on page 455 shows an IBM 7207 1/4" cartridge tape drive.



Figure 79. IBM 7207 1/4" Cartridge Tape Drive

The 7207 is the latest enhancement to the IBM 7207 1/4-inch family of tape drives. It provides an excellent, low-cost, entry based solution for tape backup with the RS/6000 family of workstations and servers. The 7207-122 incorporates the most recent Single Channel Linear Recording (SLR5) QIC technology.

Table 146 summarizes the key features of the 7207 Model 122 tape drive.

Table 146. IBM 7207-122 4 GB SLR5 Tape Drive at a Glance

Characteristics	7207 Model 122
Tape drive type	QIC (1/4-inch)
Capacity	4 GB (8 GB with compression)

Characteristics	7207 Model 122
Media	4 GB data cartridge (P/N 59H3660)
Max. data transfer rate	380 KB/s (up to 760 KB/s with compression)
Interface	SCSI-2
Supported systems	Supported on selected RS/6000 models using SCSI-2 differential interface
Operating systems supported	AIX Version 4.2, 4.3, or later

13.3.4 IBM 7208 8 mm Tape Drive

Figure 80 shows an IBM 7208 8 mm tape drive.



Figure 80. IBM 7208 8 mm Tape Drive

Improved design, along with error correction code (ECC), provides enhanced reliability and makes this drive a high-capacity, low-cost-per-megabyte storage solution for unattended stand-alone backup. An SCSI cable, 8 mm data cartridge, cleaning cartridge, and test tape are included.

7208-341

- Enhanced reliability and performance
- Up to 20 GB / 40 GB data storage capacity
- Up to 3 MB/s / 6 MB/s data transfer rate

Table 147 summarizes the key features of the 7208 Model 341 8 mm tape drive.

Table 147. IBM 7208-341 8 mm Tape Drive at a Glance

RS/6000 Hardware Requirements	SCSI-2 Differential High-Performance I/O Controller (# 2420) Enhanced SCSI-2 Differential Fast/Wide Adapter/A (# 2412) SCSI-2 Differential Fast/Wide Adapter/A (# 2416) PCI SCSI-2 Differential Fast/Wide Adapter (# 2409, # 6209) PCI Differential Ultra SCSI Adapter (# 6207)
--------------------------------------	---

Operating Systems	AIX 4.2.1, 4.3, or later
Media	59H2678
Warranty Period	One-year limited warranty
Dimensions	2.2"H x 9.8"W x 10.8"D (55 mm x 250 mm x 275 mm)

13.4 IBM Tape Automation Products

This section covers the following tape automation products available from IBM:

- 3490 E Magnetic Tape Subsystem Enhanced Capability Models
- Magstar 3590 Tape Drive
- Magstar 3494 Tape Library
- Magstar 3570 Tape Subsystem
- Magstar 3575 Tape Library Dataserver
- 7332 4 mm DDS Tape Autoloader
- 7331 8 mm Tape Library
- 7337 Digital Linear Tape Library

The section starts with notes on product positioning. A summary table for the tape automation products and more detailed descriptions of the products follow.

13.4.1 Tape Automation Product Positioning Notes

The primary tape automation products for RS/6000 systems are the Magstar products:

- Magstar 3570 Multi Purpose (MP) Tape Subsystem and the Magstar 3575 MP Tape Dataserver offer the best combination of data rate and access time among midrange and network tape storage solutions.
- Magstar 3590 tape drives are installed in IBM and non-IBM environments characterized by the need for large storage and high performance.
- Magstar 3494 tape library fulfills large tape library automation requirements.

13.4.2 Tape Automation Product Summary

Table 148 and Table 149 summarize the key features of the tape automation products available from IBM.

Table 148. IBM Tape Automation Products

Machine Model	3570 -C00 Drive	3570 -C01 Library _C11 Rack	3570 -C02 Library -C12 Rack	3575 L06, L12, L18, L24, L32	3490E -F00 Table -F01 Deskside -F11 Rack -F1A Library	3590 -B1A Library -B11 Rack	3590 -E1A Library -E11 Rack
Description	Magstar MP 3570 Drive	Magstar MP 3570 Library	Magstar MP 3570 Library	Magstar MP 3575 Library	1/2-inch Drive	1/2-inch Drive	1/2-inch Drive
Number of Drives	1	1	2	1-6	1	1	1
Max. Number of Cartridges	1	20 20	20 20	60-324	1 10 10 N/A	1 10	1 10
Native Cartridge Capacity	7 GB	7 GB	7 GB	7 GB	800 MB	20 GB ⁹	20 GB ⁹
Max. Native Capacity per Drive/Library	7 GB	140 GB 140 GB	140 GB 140 GB	420 GB to 2.2 TB	800 MB 8 GB 8 GB N/A	20 GB ⁹ 200 GB	20 GB ⁹ 200 GB
Max. Compressed¹	21 GB	420 GB	420 GB	1.2 TB 6.6 TB	2.4 GB	60 GB 600 GB	60 GB 1.2 TB
Drive Data Rate per sec.	7 MB ⁷	7 MB ⁷	7 MB ⁷	7 MB ⁷	3 MB ⁷	9 MB ⁷	14 MB ⁷
Max Compressed Data Rate¹	15 MB ⁷	15 MB ⁷	15 MB ⁷	15 MB ⁷	6.8 MB ⁷	27 MB ⁷	34 MB ⁷
Interface	SCSI-2 ⁶	SCSI-2 ⁶	SCSI-2 ⁶	SCSI-2 ⁶	SCSI-2 ⁶	SCSI-2 ⁶	SCSI-2 ⁶
Supported Platforms	IBM SUN HP-UX Windows NT	IBM SUN HP-UX Windows NT	IBM SUN HP-UX Windows NT	IBM SUN HP-UX Windows NT	IBM	IBM SGI SUN HP-UX Windows NT	IBM SGI SUN HP-UX Windows NT

Table 149. IBM Tape Automation (Continued)

Machine Model	3494 -L10/D10/S10	3494 -L12/D12/S10	7332	7331-305	7337 -305/306 (rack)
Description	3490-CxA/FIA Library	3590-B1A/E1A Library	DDS-3 Autoloader	8 mm Library	DLT Library
Number of Drives	1-16	1-16	1	1-2	1-2
Max. Number of Cartridges	6240 ⁴	6240 ⁴	4	20 ²	15
Native Cartridge Capacity	800 MB	20 GB	12 GB	20 GB	35 GB
Max. Native Capacity per Drive/Library	2.4 TB ³	124.8 TB ⁵	48 GB	440 GB	525 GB
Max. Compressed ¹	7.2 TB ³	374 TB ⁵	96 GB	880 GB	1.05 TB
Drive Data Rate per sec.	3 MB ³	9 MB ⁵ /14 MB	1.1 MB	3 MB	5 MB
Max Compressed Data Rate ¹	6.3 MB ³	13 MB ⁵ /34 MB	2.2 MB	6 MB	10 MB
Interface	SCSI-2 ⁶ ESCON OEMI	SCSI-2 ⁶	SCSI-2 F/W SE	SCSI-2 ⁶	SCSI-2
Supported Platforms	IBM SUN CRAY	IBM SGI SUN	IBM	IBM	IBM

Notes:

1. Compression rate is data dependent.
2. Plus 2 bonus slots for cleaning, test, data cartridges.
3. Capacity/data rate calculated, assuming all drives are 3490s.
4. Maximum number of cartridges decreases as tape/control units are added.
5. Capacity/data rate calculated, assuming all drives are 3590s.
6. SCSI-2 Fast/Wide Differential.
7. 3570 Cxx drives, data rate for Bxx drives is 2.2 MB/s.
8. 3570 B00, B01, B11, B02, and B12 are also available.
9. 3590 values calculated using Extended Length Cartridge Capacity.

13.4.3 IBM 3490E Magnetic Tape Subsystem Models

Figure 81 shows the IBM 3490E Magnetic Tape Subsystem.



Figure 81. IBM 3490E Magnetic Tape Subsystem

The IBM 3490 Enhanced Capability (3490E) F models are designed to provide outstanding subsystem performance, high cartridge capacity, high reliability, and operational efficiency while facilitating remote tape operations.

Using a 36-track head while reading or writing 18 tracks bidirectionally, the 3490E F Models double the capacity of a base IBM Cartridge System Tape and virtually eliminate rewind time for a full cartridge. The 3490E F Models also are read-compatible with 3480-track tape. The RS/6000 and AS/400 exploit the new 3480 read interchange non-standard write function.

Major design enhancements for the 3490E F models improve mechanical and data reliability. Tape path enhancements, advanced electronic packaging, and circuit technology, along with enhanced control unit error recovery procedures, make the 3490E F models extremely reliable.

The Model F00 is a low-priced tabletop model. The Model F01 is a desk-side tape library with a 10-cartridge auto loader. The Model F11 has a 10 cartridge auto loader and is designed for rack-mounting. The Model F1A is designed for installation in the Magstar 3494 Tape Library.

They are smaller and lighter-weight full-function versions of the 3490 36-track, 1/2-inch cartridge tape family. They have the following attributes:

- Up to 6.8 MB/s
- Up to 2.4 GB per cartridge
- Up to 24 GB random access loader (10 cartridges)

- Reads and writes 18-track tape cartridges and reads and writes 36-track tape cartridges
- Self-functioning cleaning cartridge

Table 150 lists the 3490E Model F01 and F11 features.

Table 150. IBM 3490E Model F01/F11 Magnetic Tape Subsystem at a Glance

Max. Tape Drives	1
Max. Cartridges	10
Capacity	0.8 GB/2.4 GB per cartridge
RS/6000 Hardware Requirements	SCSI-2 Fast/Wide Differential Adapter/A (# 2416) SCSI-2 Differential High-Performance I/O Controller (# 2420); Enhanced SCSI-2 Differential Fast/Wide Adapter/A (# 2412); PCI SCSI-2 Differential Fast/Wide Adapter (# 2409) PCI SCSI-2 Fast/Wide Adapter (# 6209)
Operating Systems	AIX Version 3.2.0+ with # 9603, the 3490 Enhanced Capability F Models; AIX Device Driver and AIX Version 4.1+
Media	09G4494
Warranty Period	One-year limited warranty
Dimensions	Model F01 Stand-alone 19.1" H x 7" W x 24.1" D (48.6 cm x 17.7 cm x 61.2 cm)
	Model F11 Rack-mount 6.8" H x 17" W x 24" D (17.1 cm x 43.2 cm x 61.1 cm)
	Model F00 Tabletop 10.2" H x 8.75" W x 17" D (25.9 cm x 22.2 cm x 43.2 cm)
	Model F1A Library 17" H x 6.8" W x 24" D (43.2 cm x 17.1 cm x 61.1 cm)

You can install one or two F1A drives in any new 3494 L10 or D10 frame on independent sliders that enable the drives to be serviced from the rear. This prevents disruption to the 3494 service when the drive is being serviced. (Note that the F1A cannot be installed in existing L10 frames due to the current frame design. The autoloader feature is not available on the F1A).

The F1A can attach to the RS/6000 or SP through ESCON or SCSI.

13.4.4 IBM Magstar 3590 Tape Subsystem

Figure 82 shows the IBM Magstar 3590 Tape Subsystem in an IBM rack.



Figure 82. IBM Magstar 3590 Tape Subsystem

The challenge in today's computing environments is to find an open, multiplatform storage solution that is both high-performance and low cost. An industry standard with over 30,000 Magstar 3590 tape drives installed, the 3590 provides the industry's highest combination of capacity and performance among tape drives designed to handle a broad spectrum of applications. Capacity is 50 times that of the 3480 and three times its data transfer rate. The Magstar 3590 is designed for 100 times the data integrity of IBM 3480 tape drives. Even large storage, high-demand online applications count on Magstar to deliver data quickly, reliably, and inexpensively.

IBM Magstar 3590 Tape Subsystem

Whether you want to create tapes for archive files, restore files, enjoy low cost per megabyte, or step up from your current 3490E or Magstar 3494 Tape Library configuration, the Magstar 3590 Tape Subsystem is the hottest solution in its class.

The Magstar 3590 is available in five models for SCSI and UltraSCSI attachment:

- The Magstar 3590 Models B11 and E11 are rack-mounted and incorporate a 10-cartridge ACF for high-capacity unattended operation. The Models B11 and E11 can be modified to Models B1A or E1A.

- The Magstar 3590 Models B1A and E1A have no ACF and are designed to be incorporated into the Magstar 3494 tape library.
- The Magstar Model C12 frame, with one to four Model B1A or E1A tape drives, provides attachment to the StorageTek 4410 and 9310 ACS.

The IBM Magstar 3590 Tape Subsystem provides high levels of capacity, performance, and reliability for S/390, RS/6000, AS/400, and a variety non-IBM UNIX and Windows NT systems. It offers significant enhancements over IBM 3480, 3490, or 3490E tape drives and demonstrates IBMs continued leadership in storage solutions. The 3590 subsystem includes tape drives and controllers that can be configured in a broad range of solutions that include stand-alone frames or racks, or in automated library solutions to meet a diverse set of customer requirements. The Magstar 3590 Tape Drives can be attached to a SCSI host interface of selected IBM and non-IBM midrange and open systems. The tape drives can also attach to ESCON channels on S/390 systems through a Magstar 3590 controller.

Model 3590-A14

The Model A14 is a frame box that can be attached to a Magstar 3495 Tape Library or used for non-library attachment of the 3590 to ES/9000, ES/3090, or S/390 systems. Four Model B1A drives, along with one Model A50 or A00 controller, can be installed in the Model A14 for attachment to a Magstar 3495 Tape Library, or up to four Model B11 drives and one Model A60 Controller (or up to two Model A50 or A00 Controllers), or up to four Model E11 drives and one Model A60 Controller (or up to two Model A50 Controllers), can be installed in the Model A14 for non-library usage on an ES/9000, ES/3090, or S/390 system.

Model 3590-A50

The 3590 Model A50 is a controller designed for attaching 3590 drives to ESCON channels on RS/6000, ES/3090, ES/9000, or S/390 systems. Attachment to RS/6000 servers is supported by using the PCI ESCON adapter (# 2751) with the PCI ESCON Channel Tape Attach AIX LPP (5765-E04). It can be installed in a Magstar 3494 Tape Library, 3590 Model A14 Frame, 3590 Model C14 Silo Compatible Frame, or standard 19" rack (such as the IBM 9309 Rack Enclosure Model 002). It can also be attached as a free-standing unit to a 3590 Model C12 Silo Compatible Frame. Its performance is roughly double what was achieved in the predecessor Model A00 Controller.

Model 3590-A60

The 3590 Model A60 is a controller designed for attaching 3590 drives to ESCON channels on RS/6000, ES/3090, ES/9000, or S/390 systems.

Attachment to RS/6000 servers is supported by using the PCI ESCON adapter (# 2751) with the PCI ESCON Channel Tape Attach AIX LPP (5765-E04). The list of supported tape devices now includes the Magstar 3590 A60 Tape Controller in native or 3490E emulation mode with Magstar 3590 E11, E1A, B11, or B1A Tape Drives. In addition, the Magstar 3590 A50 Tape Controller with 3590 E11 or E1A drives is now supported. It can be installed in a Magstar 3494 Tape Library, 3590 Model A14 Frame, 3590 Model C10 Silo Compatible Frame, or standard 19" rack (such as the IBM 9309 Rack Enclosure Model 002, 7014 RS/6000 Rack Model S00, or the 7015 RS/6000 System Rack Model R00). Its offers significant performance improvements over the Model A50 Controller.

Model 3590-B11

The Magstar 3590 Model B11 tape drive is the rack-mountable model that includes a 10-cartridge Automatic Cartridge Facility (ACF) that can be used in a random access mode as a mini-library. Up to four Model B11s can be installed in a rack, thus reducing the floor space required. The ACF can be quickly loaded with a new magazine. This model includes an integrated SCSI-3 controller with two ports that support a 16-bit, fast-and-wide, Ultra SCSI differential interface. It can also be mounted in a 3590 Model A14 Frame with the 3590 Model A60, A50, or A00 Controller for non-library attachment to ESCON channels. The Model B11 tape drive has a 9 MB/s native data rate (with compression up to 27 MB/s sustained data rate).

Model 3590-B1A

The Magstar 3590 Model B1A tape drive comes without the ACF and is designed to go into the Magstar IBM 3494 Tape Library. It can also go in the IBM 3590 Model C12 or C14 Silo Compatible Frame for attachment to a StorageTek Automated Cartridge System (ACS) or in the 3590 Model A14 Frame for attachment to a Magstar 3495 Tape Library. This model includes an integrated SCSI-3 controller with two ports that support a 16-bit, fast-and-wide, Ultra SCSI differential interface. The Model B1A tape drive has a 9 MB/s native data rate (with compression up to 27 MB/s sustained data rate).

Model 3590-E11

The Magstar 3590 Model E11 tape drive is the rack-mountable model that includes a 10-cartridge Automatic Cartridge Facility (ACF) that can be used in a random access mode as a mini-library. Up to four Model E11s can be installed in a rack, thus reducing the floor space required. The ACF can be quickly loaded with a new magazine. This model includes an integrated SCSI-3 controller with two ports that support a 16-bit, fast-and-wide, Ultra SCSI differential interface. It can also be mounted in a 3590 Model A14 Frame with the 3590 Model A60 or A50 Controller for non-library attachment

to ESCON channels. The Model E11 tape drive has a 14 MB/s native data rate (with compression up to 34 MB/s sustained data rate).

Model 3590-E1A

The Magstar 3590 Model E1A tape drive comes without the ACF and is designed to go into the Magstar IBM 3494 Tape Library. It can also go in the IBM 3590 Model C12 or C14 Silo Compatible Frame for attachment to a StorageTek Automated Cartridge System (ACS). This model includes an integrated SCSI-3 controller with two ports that support a 16-bit, fast-and-wide, Ultra SCSI differential interface. The Model E1A tape drive has a 14 MB/s native data rate (with compression up to 34 MB/s sustained data rate).

Existing Magstar Model B11 and B1A Tape Drives (B Models) can be field-upgraded to the new E Models, which helps protect existing investments in Magstar 3590 tape drives. Media investments are also protected because the new E Model drives can both read (128 track) and write (256 track) to existing cartridges. IBM will continue to offer Magstar 3590 B Models with the recently enhanced Ultra SCSI support along with the new E models.

13.4.5 IBM Magstar 3494 Tape Library

Figure 83 shows the IBM Magstar 3494 Tape Library.



Figure 83. IBM Magstar 3494 Tape Library

The Magstar 3494 Tape Library is a tape automation system that consists of individual frame units for modular expansion that provides a wide range of

configurations. This flexibility enables organizations to start small and grow in an affordable and incremental manner.

The basic building block of the Magstar 3494 is the Lxx Control Unit Frame, which contains a library manager, a cartridge accessor, up to two tape drives, and slots for the storage of tape cartridges. To the Lxx frame, you can add drive frames, storage unit frames, the Magstar Virtual Tape Server, and a high-availability model to create a maximum configuration of 16 frames and two service bays.

The Magstar 3494 supports both Magstar 3590 Model B1A/E1A and 3490E Model F1A tape drives within the same library but in separate frames.

Highlights:

- Provides cost-effective, reliable, and space-efficient tape automation for a variety of environments.
- Uses both Magstar* 3590 Model B1A/E1A and 3490E tape drives.
- Supports the new Magstar 3590 Model A60 High-Performance ESCON* Control Unit.
- Provides a high-availability model with an optional feature that enables two active accessors for increased performance.
- Provides multiple host connectivity options.
- Offers exceptional expandability from one to 16 frame units from one to 76 Magstar 3590 (with ESCON) or 32 3490E tape transports and from 160 to 6,240 cartridges (up to 748 TB(1) with 3:1 compression).
- Can mount/demount up to 610 cartridges per hour with dual active accessors(2).
- With the Magstar Virtual Tape Server, you can utilize the full capacity of Magstar 3590 tape cartridges to significantly reduce overall tape operating costs.
- Supports the IBM 3466 Network Storage Manager.

The average mount-access time in a single-frame Magstar 3494 is only seven seconds. The Magstar 3494 can perform up to 265 cartridge exchanges per hour with a single gripper, up to 305 exchanges per hour with the optional dual gripper, and up to 610 exchanges per hour with a dual gripper and dual active accessors. The exchange capability of the accessor can decrease as the number of frames increases because of increased accessor travel.

Magstar 3494 library components consists of a control unit frame, a drive unit frame, and a storage unit frame.

The control unit frame provides the library manager, console, cartridge accessor, optional convenience I/O stations, optional tape drive, barcode reader, optional ESCON controller, and up to 240 cartridge storage cells. Configurations include:

- Model L10--One or two 3490E Model F1A tape drives.
- Model L12--Zero to two SCSI-attached Magstar 3590 tape drives.
- Model L14--Zero to two Magstar 3590 tape drives and an ESCON control unit (3590-A50).

The drive unit frame has the following characteristics:

- The Magstar 3494 drive unit frame provides the capability of adding tape drives and up to 400 cartridge storage cells. Configurations include:
 - Model D10--Zero to two 3490E Model F1A tape drives.
 - Model D12--Zero to six SCSI-attached Magstar 3590 tape drives.
 - Model D14--Zero to four Magstar 3590 tape drives and an ESCON control unit 3590-A50 or 3590-A60).
- A Magstar 3590 A60 control unit installed in a D14 drive unit frame can provide ESCON support for 3590 tape drives in an adjacent L12, L14, or D12 library frame.

The storage unit frame has the following characteristics:

- The Magstar 3494 storage unit frame provides the capability for adding cartridge storage.
 - Model S10--up to 400 cartridge cells.
 - Magstar Virtual Tape Server-3494 Model B18 Frame
- The Magstar Virtual Tape Server contains a RISC-based microprocessor, storage management functions and up to 864 GB of disk cache (3:1 compression).
- The Extended High-Performance Option enables data compression for host channel attachment and provides larger effective disk capacities, improved performance, and 64 virtual device addressing. Up to 300,000 virtual volumes can be stored in a Magstar 3494 Tape Library with two Model B18 units.

IBM Magstar 3494 Tape Library		
Media	05H4434 with 3590 09G4494 with 3490E	
Warranty Period	One-year limited warranty	
Dimensions	Model L10/12/14	Control Unit 70.9"H x 29.5"W x 60.0"D (180 mm x 75 mm x 152.5 mm)
	Model D10/12/14	Drive Unit 70.9"H x 29.5"W x 60.0"D (180 mm x 75 mm x 152.5 mm)
	Model S10	Storage Unit 29.3"H x 29.5"W x 29.3"D (180 mm x 75 mm x 75 mm)
Notes: (1) Assumes maximum compression ratios. Compression ratios can vary widely based on data characteristics. (2) Does not support attachment to IBM 3590.		

13.4.6 IBM Magstar MP 3570 Tape Subsystem

Figure 84 shows the IBM Magstar MP 3570 Tape Subsystem.



Figure 84. IBM Magstar MP 3570 Tape Subsystem

The Magstar MP significantly expands the functional utility of tape storage with its revolutionary data recall performance, rugged cartridge designed for automation, and super-reliable Magstar technology.

Magstar MP is also well suited for traditional tape backup applications—offering improved price/performance, reliability, and capacity over traditional mid-range technologies. Magstar MP will change the way you think about tape storage.

- IBM-patented, midpoint load tape cartridge, ready to fast forward to your data
- Unique tape path eliminates tape thread time for reduced complexity and higher reliability. Quick-load tape never leaves the cartridge
- Backed by an industry-leading, three-year warranty
- Offering one- or two-drive library configurations
- Self-contained tape cleaning system with an internal cartridge slot
- Total library capacity up to 420 GB¹
- Supported by leading storage management applications, including ADSM for AIX, Sysback for AIX, BRMS for AS/400
- 2.2 MB/s to 15 MB/s¹ data transfer rate
- Native device support for Sun, Windows NT, and HP-UX
- New 3570 C models provide up to 50 GB/hour¹ data backup per drive
- Optional bar code reader for B models
- Standard bar code reader for C models

Table 152 summarizes the key features of the Magstar 3570 Tape Subsystem.

Table 152. IBM Magstar MP 3570 Tape Subsystem at a Glance

Model	B00/C00	B01/C01	B11/C11	B02/C02	B12/C12
Max. Tape Drives	1	1	1	2	2
Max. Cartridges	1	20	20	20	20
Packaging	External Drive	Tabletop Library	Rack-mount Library	Tabletop Library	Rack-mount Library
Capacity	7 GB/21 GB ¹ per cartridge				
Characteristics	3:1 LZ1 compression Recording technique No. of tracks		Standard Interleaved serpentine longitudinal 128		
RS/6000 Hardware Requirements	SCSI-2 Fast/Wide Differential Adapter/A (# 2416); SCSI-2 Differential High-Performance I/O Controller (# 2420); Enhanced SCSI-2 Differential Fast/Wide Adapter/A (# 2412); PCI SCSI-2 Differential Fast/Wide Adapter (# 2409, # 6209)				
Operating Systems	# 2416, # 2420, # 2412: AIX Version 3.2.5 and AIX Version 4.1.1+ # 2409, # 6209: AIX Version 4.1.1+				
Device Driver	AIX 3.2.5, AIX 4.1.1+ provided. Unattended backup requires use of programs such as ADSM, other tape management software, or user-defined scripts utilizing robotics commands.				
Media	CSH2462				
Warranty Period	Three-year limited warranty				
Dimensions					
Model	B00 C00	B01/B02 C01/C02		B11/B12 C11/C12	
Height	4.4" 112 mm	8.5" 217 mm		8.5" 217 mm	
Width	12.6" 320 mm	19.0" 438 mm		17.5" 444 mm	
Depth	13.3" 338 mm	30.4" 771 mm		28.1" 714 mm	
EIA Units				5	
Ordering Highlights					
The library models include two removable 10-cartridge magazines, cleaning cartridge, blank data cartridge, and power cord. At initial order only, choose # 9570 1 cleaning/1 data cartridges or # 8701 1 cleaning/10 (price of 9 at a lower price) data cartridges or # 8702 1 cleaning/20 data (price of 19 at a lower price) cartridges. Each drive has two SCSI controllers. B00 cover is black. B01 and B02 are either white or black. B11/B12 are white.					
¹ Assumes maximum compression ratios. Compression ratios can vary widely based on data characteristics.					

Model	B00/C00	B01/C01	B11/C11	B02/C02	B12/C12
F/C	Description				
2891	Interposer HD50/HD68 for # 2420				
2892	Interposer CC68/HD68 for # 2412, # 2416				
5205	0.5 M SCSI cable				
5245	4.5 M SCSI cable				
5212	12 M SCSI cable				
5218	18 M SCSI cable				
8701	10 data cartridges/1 cleaning cartridge				
8702	20 data cartridges/1 cleaning cartridge				
8752	10 data cartridges				
9570	1 data cartridge/1 cleaning cartridge				
9600	Attach to RS/6000				
9076	Attach to RS/6000 SP				
¹ Assumes maximum compression ratios. Compression ratios can vary widely based on data characteristics.					

13.4.7 IBM Magstar MP 3575 Tape Library Dataserver

Figure 85 shows the IBM Magstar MP 3575 Tape Library Dataserver.



Figure 85. IBM Magstar 3575 Tape Library Dataserver

The IBM Magstar MP 3575 Tape Library Dataserver is a new family of automated storage solutions designed for the growing, unattended storage requirements of today's midrange systems and network servers. These compact, integrated tape storage libraries expand the capability of tape processing by optimizing both read- and write-intensive operations. A dual-gripper picker can provide fast cartridge exchange times between the library slots and the Magstar MP tape drives in the library. The Magstar MP 3575 Tape Library attaches to RS/6000 and other SCSI-attached open systems in single or multi-host configurations.

- Offers five new models for the SCSI open systems environments.
- Includes Magstar MP tape drives that provide fast data access for current and emerging applications, such as save/restore, network storage management, data warehousing, and digital libraries.
- Increases the amount of data that can be accessed with near online performance of up to 6.8 TB of storage capacity (with 3:1 compression).
- Delivers a maximum sustained data transfer performance of 50 to 300 GB/hour combined throughput (with 3:1 compression).
- Provides multi-host attachment with up to three logical libraries.
- Supports industry-leading storage management offerings to provide enterprise-wide backup/restore and archive/retrieval.
- Supports Magstar 3570 MP C models.
- Supports AS/400, AIX, HP-UX, Sun, and Windows NT.

Table 153 summarizes the key features of the Magstar 3575 Tape Subsystem.

Table 153. IBM Magstar MP 3575 Tape Library Dataserver at a Glance

Model	L06	L12	L18	L24	L32
Capacity (native) (compressed)	420 GB 1.2 TB	840 GB 2.5 TB	1.2 TB 3.7 TB	1.6 TB 5.0 TB	2.2 TB 6.8 TB
Cartridges	60	120	180	240	324
Drives	2	2-4	2-6	2-6	2-6
Average Move Time	<4 sec	<4 sec	<4 sec	<4 sec	<4 sec
Dimensions					
Height	39.0" 991 mm	40.5" 1029 mm	40.5" 1029 mm	50.75" 1518 mm	50.75" 1518 mm
Width	14.0" 355 mm	39.7" 1009 mm	39.7" 1009 mm	39.7" 1009 mm	39.7" 1009 mm
Depth	32.9" 836 mm	33.9" 861 mm	33.9" 861 mm	33.9" 861 mm	33.9" 861 mm
Weight	157 lbs 71 kg	280 lbs 127 kg	290 lbs 132 kg	428 lbs 195 kg	446 lbs 203 kg
Drive Characteristics	LZ1 compression Recording technique No. of tracks SCSI-2 Interface Capacity per cartridge		Standard Interleaved serpentine longitudinal 128 Fast/Wide Differential 5 GB up to 15 GB (with 3:1 LZ1 compression)		
Drive Performance	Native sustained data rate Compressed sustained data rate Drive load-to-ready time Drive average search time		7 MB/s ¹ 21 MB/s (with 3:1 compression) 6.7 sec 8.0 sec		

Model	L06	L12	L18	L24	L32
Operating Environment	Temperature Relative humidity Wet bulb Heat output		60°-90° F (16°-32°C) 8% to 80% non-condensive 78° F (26° C) maximum		
	Noise levels (operating)		L06: 600 BTU/hr L12: 850 BTU/hr L18/L24/L32: 1200 BTU/hr		
	Power requirements		L06: 52 dBs L12: 53 dBs L18/L24/L32: 54 dBs		
Standard Features for all Models	100-127 VAC or 200-340 VA (with 3:1 LZ1 compression)				
Software Support	Auto cleaning, bar code reader, bulk load door, dual gripper, I/O station, keylock, library viewing window, operator control panel, and self-contained diagnostic				
Media	ADSM V2.1.0.10 and V2.1.5.10 NetTAPE for AIX AIX 3.2.5.4.1.x and later, and 4.2.0 and later CA/Cheyenne ARCserve 6.0 Legato Systems NetWorker 5.0 Spectra Logic Alexandria 3.71/3.75 SCH Technologies REELlibrarian and REELback Veritas				
	05H2462				
¹ With 3570 C model drives.					

13.4.8 IBM 7332 4 mm DDS Tape Autoloader

Figure 86 shows the IBM 7332 Model 220 4 mm Tape Autoloader.



Figure 86. IBM 7332 Model 220 4 mm Tape Autoloader

The demand for dependable, cost-effective tape automation solutions exists in all data processing environments, including RS/6000 systems. Increases in data storage to improve customer service and provide better business solutions require fast, dependable, and affordable disk-to-tape backup operations and large data archive and retrieval facilities.

13.4.8.1 IBM 7332 Model 110

The IBM 7332 Model 110 4 mm DDS-3 Tape Cartridge Autoloader extends the IBM family of 4 mm tape offerings. It now has an enhanced autoloader for attachment to the RS/6000 family of workstations and servers. The 7332 Model 110 is a multiple cartridge tape autoloader with one 4 mm tape drive that operates in streaming mode and uses Digital Audio Tape (DAT) technology. This externally attached tape autoloader attaches to a single-ended SCSI adapter via SCSI cable and provides additional tape storage capacity for save and restore operations.

The 7332 Model 110 uses 12 GB capacity cartridges to obtain 48 GB capacity with the four cartridge magazine. With 2:1 compression, the 7332 tape autoloader magazine can contain up to 96 GB capacity. The 4 mm tape drive in the 7332 Model 110 has a sustained data rate of up to 1.1 MB/s (2.2 MB/s with compression). It contains a media recognition feature for greater reliability by allowing only data grade tape cartridges to be used for writing data. The 7332 tape autoloader incorporates state-of-the-art DAT technology, which uses Digital Data Storage (DDS) tape media. In addition to reading and writing on DDS-3 formatted tape cartridges, the 7332 Model 110 also reads and writes on DDS or DDS-2 tape cartridges.

13.4.8.2 IBM 7332 Model 220

The IBM 7332 Model 220 4 mm DDS-4 Tape Cartridge Autoloader is a multiple cartridge tape autoloader with one 4 mm DDS-4 tape drive. It features data compression that provides an effective capacity of up to 240 GB (with 2:1 compression) and a sustained data rate of up to 3.0 MB/s (6.0 MB/s with compression), depending on the characteristics of the data transferred. User productivity is enhanced with a three times larger tape cartridge capacity and higher data rate over the previous 7332 Model 110. The 7332 Model 220 provides a more cost-effective solution for save and restore and archiving functions than previous models.

The IBM 7332 Model 220 4 mm DDS-4 Tape Cartridge Autoloader provides an integrated, cost-effective, entry-level tape storage solution. It is based on proven 4 mm tape drive technology and a tape autoloader.

Note

To connect a #6205 to the following devices:

- 7204-409 External 9.1 GB disk drive
- 7204-419 External 18.2 GB disk drive
- 7206-220 External 4 mm tape drive
- 7332-220 External 4 mm tape drive

you need to order no-charge VHDCI Cable/Interposer (# 9799) and one of the appropriate cables:

- 9137 1.0 m Ultra2 SCSI cable
- 9149 1.0 m SCSI-2 F/W cable

13.4.9 IBM 7331 8 mm Tape Library

Figure 87 shows the IBM 7331 8 mm Tape Library.



Figure 87. IBM 7331 8 mm Tape Library

This high-quality tape library has the capacity of similarly priced products and automatic drive cleaning. The dual-host feature allows two host processors to be attached to a 7331 Tape Library, equipped with two tape drives, and operate as two independent libraries.

Whether you need save/restore capabilities, distribution, migration, or mass archiving, it's the ideal low-cost, high-performance solution. And, it is IBM designed and built with your high-production environment in mind.

- Enhanced “Mammoth” technology

- Exceptional value in this IBM-designed and developed tape library
- Total capacity up to 20 cartridges (40 GB with compression)
- 440 GB library capacity (880 GB with compression)
- Up to 3 MB/s to 6 MB/s data transfer rate
- Supports one or two 20 GB tape drives (2nd drive is field installable)
- Supported by leading storage management applications including ADSM for AIX
- Compatible with industry-leading backup software

Table 154 summarizes the key features of the 7331 Model 305 8 mm Tape Library.

Table 154. IBM 7331 Model 305 8 mm Tape Library at a Glance

Max. Tape Drives	2
Max. Cartridges	20 plus 2 (cleaning tape and one user defined purpose)
Capacity	20 GB/40 GB ¹ per cartridge (P/N 59H2678)
RS/6000 Hardware Requirements	SCSI-2 Differential High-Performance External I/O Controller (# 2420); SCSI-2 Differential Fast/Wide Adapter/A (# 2416); Enhanced SCSI-2 Differential Fast/Wide Adapter/A (# 2412); PCI SCSI-2 Differential Fast/Wide Adapter (# 2409, # 6209); PCI Differential Ultra SCSI Adapter (# 6207); HP 9000 Series 700 and 800 models supporting EISA bus; Selected Sun workstations
Operating Systems	AIX 3.2.5, 4.1.1, 4.1.5 or 4.2.0 and later releases HP-UX Solaris 2.3, 2.4 or 2.5
Media	59H2678
Warranty Period	One-year limited warranty
Dimensions	25.1" H x 12.7" W x 28.5" D (637 mm x 323 mm x 723 mm)
¹ Assumes maximum compression ratios. Compression ratios can vary, widely based on data characteristics.	

13.4.10 IBM 7337 Digital Linear Tape Library

Figure 88 on page 479 shows the IBM 7337 Digital Linear Tape Library.



Figure 88. IBM 7337 Digital Linear Tape Library

The IBM 7337 provides an integrated, cost-effective, high-function, automated tape library solution for selected RS/6000 models. It is based on proven DLT drive technology and industry-leading library automation. Its innovative design incorporates 15 data cartridges and one or two DLT7000 drives in a rack-mount or desktop configuration that minimizes the library's physical dimensions while maximizing available capacity.

The IBM 7337 is designed for capacity, speed, reliability, and ease of use and is backed by worldwide service and support. For expanded library and media management, the 7337 can be used with IBM ADSM or other industry-leading storage management software.

- Affordable, automated, digital linear tape library
- Total library capacity up to 15 cartridges or 1.05 TB (compressed data)
- Supports one or two 35 GB tape drives (2nd drive is field installable)
- 35 GB cartridge capacity (70 GB with compression)
- Up to 5 MB/s to 10 MB/s data transfer rate
- Bar code reader
- Supported on AIX 4.1.5, AIX 4.2.0, AIX 4.2.1 and later releases
- Supported by leading storage management applications, including ADSM for AIX
- Read and write cartridges compatible with IBM 7205 Model 311
- Read and write compatible with previous DLT media formats
- Model 305—Desktop, Model 306—Rack

Table 155 summarizes the key features of the 7337 Digital Linear Tape Library.

Table 155. IBM 7337 Digital Linear Tape Library at a Glance

Max. Tape Drives	2
Max. Cartridges	15 plus
Capacity	35 GB/70 GB ¹ per cartridge
RS/6000 Hardware Requirements	PCI SCSI-2 Differential Fast/Wide Adapter (# 2409, # 6209) Enhanced SCSI-2 Differential Fast/Wide Adapter/A (# 2412) SCSI-2 Differential Fast/Wide Adapter/A (# 2416); SCSI-2 Differential High-Performance I/O Controller (# 2420) PCI Differential Ultra SCSI Adapter (# 6207)
Operating Systems	AIX Version 4.1.5; AIX 4.2.0; AIX 4.2.1 and later
Media	59H3040
Warranty Period	One-year limited warranty
Dimensions	Model 305 9.25" H x 18.9" W x 26.5" D (235 mm x 480 mm x 673 mm) Model 306 8.75" H x 17.5" W x 26.5" D (222 mm x 445 mm x 673 mm)
¹ Assumes maximum compression ratios. Compression ratios can vary widely based on data characteristics.	

13.5 IBM External Optical Devices

The following section describes the following external optical devices available from IBM:

- 7210 CD-ROM Drive
- 3995 Optical Library
- 7209 Extended Multifunction Optical Drive

13.5.1 IBM 7210 CD-ROM Drive

Figure 89 on page 481 shows the IBM 7210 CD-ROM Drive.



Figure 89. IBM 7210 CD-ROM Drive

This compact, self-powered, read-only optical drive uses a standard CD-ROM disc to support random access online files.

The 7210 Model 020 is a high-performance CD-ROM drive from IBM's StorageSmart family of products. The Model 020 features reduced access time and four times the average data rate of previous models. Model 020 has average access time from 90 milliseconds to 175 milliseconds and data transfer rates from 150 KB/s to 5.1 MB/s. Software, technical manuals, updates, images, and other vital material on CD-ROM can be accessed quickly and efficiently.

The Model 020 supports synchronous and asynchronous modes, accommodates both 12 cm and 8 cm disks, and can read CD WORM disks created in either single- or multi-session by CD-R devices. Its downward compatibility and low cost protect your investment.

Highlights include the following:

- Reading 650 MB capacity CD-ROM disks supported
- 256 KB memory buffer included
- Attachment via industry-standard SCSI-2
- Slim 5.25-inch, half-high form-factor
- Fully multi-media (MPC-3) compliant
- Reliable and affordable 14X-32X variable speed CD-ROM drive
- Improved performance with up to 5.1MB/s data rate - four times the rate of the current 7210 Model 015
- Backward compatibility with 2X,4X, 8X, and 14X media
- Support for multi-session CD-R disks
- Reads CD-RW discs
- CD-DA standard Red Book Audio
- CD-ROM data in modes 1 and 2

- CD-ROM XA data in either Mode 2 Form 1 or Mode 2 Form 2

Table 156 summarizes the key features of the 7210 Model 020 External CD-ROM Drive.

Table 156. IBM 7210 Model 015 External CD-ROM Drive at a Glance

Media	80 mm and 120 mm disk formats supported
Access Time	150 ms typical
Speed	Runs at 1x (150 KB/s); 2x (300 KB/s); 4x (600 KB/s) 8x (1200 KB/s)
File Format	Read - ISO 9660/High Sierra file format
RS/6000 Hardware Requirements	SCSI-2 High-Performance External Controller (# 2410) SCSI-2 Fast/Wide Adapter/A (# 2415) SCSI-2 High-Performance Internal I/O Controller (# 2831) SCSI High-Performance External Controller (# 2835) PCI SCSI-2 SE Fast/Wide Adapter (# 2408, # 6208)
Operating Systems	AIX Version 3.2.5 (with PTFs), 4.1.4+, 4.2+ Microsoft Windows NT Version 3.51 and 4.0
Warranty Period	One-year limited warranty
Dimensions	2.165" H x 9.842" W x 10.826" D (55 mm x 250 mm x 275 mm)

13.5.2 IBM 3995 Optical Library Solutions: The Enhanced C-Series

Figure 90 on page 482 shows the IBM 3995 Optical Libraries.



Figure 90. IBM 3995 Optical Libraries

With low cost, high-capacity, and direct access, optical storage is an excellent solution for bringing new data and information online or for keeping existing data online for longer periods of time. New data and information includes information that is currently stored on paper or microfiche. Extended Multi-function optical drives, an IBM first-in-the-industry, support Permanent WORM, rewritable, and Continuous Composite Write Once (CCW WORM) optical cartridges—delivering up to 5.2 GB of storage on each optical cartridge.

IBM 3995 Enhanced C-Series for Open Systems Optical Libraries provide:

- Rapid access to document images, such as medical records, insurance claims, accounts payable, bills of lading, to name a few
- Extended online data life in your storage hierarchy
- Long-term media archive life to satisfy business or legal requirements for data retention
- WORM and rewritable technology within the same library to address both unalterable data storage requirements and flexible data storage needs
- Five models delivering from 104 GB to 1.3 TB of online optical storage
- Industry standard 5.2 GB, 5.25-inch optical drives for increased capacity and enhanced performance

IBM 3995 Enhanced C-Series Optical Libraries enable:

- Document imaging applications
- Computer Output to Laser Disk (COLD) applications allowing you to replace microfiche
- Report management applications, such as billing, payroll, customer lists, and accounting reports
- Data archive

Table 157 summarizes the key features of the 3995 Optical SCSI Models.

Table 157. IBM 3995 Optical SCSI Models at a Glance

Model	C60	C62	C64	C66	C68
Capacity	104 GB	270 GB	540 GB	811 GB	1341 GB
Max. Drives	1 or 2	2	2 or 4	4 or 6	4 or 6
Gripper Type	Single	Dual	Dual	Dual	Dual
Average Cartridge Move Time (sec.)	2.8	2.6	2.8	2.8	3.1

Model	C60	C62	C64	C66	C68
RS/6000 Hardware Requirements	One of the following: Single-Ended SCSI Adapters (one library per dedicated adapter); SCSI-2 High-Performance External Controller (FC #2410 or #2831); SCSI-2 Fast/Wide Adapter/A (# 2408, # 2414, # 2415, or # 6208); Differential SCSI Adapter; SCSI-2 Differential High-Performance External Controller (# 2420); SCSI-2 Differential Fast/Wide Adapter/A (# 2409, # 2412, # 2413, # 2416, or # 6209)				
Operating Systems	AIX/6000: Single-ended SCSI requires Version 3.2.0+; Differential SCSI requires Version 3.2.4+; SCSI-2 Fast/Wide Adapter/A (single-ended or differential) requires Version 3.2.5+; Selected application or optical library manager may have its own properties prerequisites. For other Operating Systems (for example Windows NT), check your application or optical library manager prerequisites.				
Other Hardware and Software Requirements	SCSI-2 Single-Ended or Differential. Library management application required to operate the 3995 (available in ADSM or from third-party providers)				
Upgrade Available	The 3995 C64 can be field upgraded to the Model C66.				
Warranty Period	One-year limited warranty				
Dimensions					
Model	C60	C62	C64	C66	C68
Height	18.0"	39.0"	40.5"	40.5"	58.25"
	457 mm	991 mm	1029 mm	1029 mm	1480 mm
Width	8.5"	14.0"	32.0"	32.0"	32.0"
	216 mm	355 mm	813 mm	813 mm	813 mm
Depth	29.0"	29.0"	30.0"	30.0"	30.0"
	737 mm	737 mm	762 mm	762 mm	762 mm
Ordering Highlights					
An SCSI cable can be shipped with each 3995. A feature code must be selected to indicate a single-ended or differential interface. The differential SCSI-2 interface provides for added distance capability and the capability to daisy-chain Model C60s and Model C62s on a single adapter.					

Model	C60	C62	C64	C66	C68
F/C	Description				
7200	Single-ended SCSI-2 Interface				
7201	Differential SCSI-2 Interface				
7210	Single-Ended SCSI-2 cable for # 2410 or # 2831 Adapter				
7211	Single-Ended SCSI-2 cable for # 2414 or # 2415 Adapter				
7212	Single-Ended SCSI-2 cable for # 2408, # 6208 Adapter				
7215	Differential SCSI-2 cable for # 2420 Adapter				
7216	Differential Daisy Chain Cable				
7217	Differential SCSI-2 cable for # 2412, # 2413, # 2416 Adapter				
7218	Differential SCSI-2 cable for # 2409, # 6209 Adapter				
7219	Single-Ended SCSI-2 cable interface for Adaptec Adapter				
7220	Differential SCSI-2 cable interface for Adaptec Adapter				
F/C	Optical Media				
8002	10 pack, Rewritable 5.2 GB, 2048 sector size				
8102	52 pack, Rewritable 5.2 GB, 2048 sector size				
8001	10 pack, Rewritable 4.8 GB, 1024 sector size				
8101	52 pack, Rewritable 4.8 GB, 1024 sector size				
8495	10 pack, Rewritable 2.6 GB, 1024 sector size				
8595	52 pack, Rewritable 2.6 GB, 1024 sector size				
8005	10 pack, CCW WORM 5.2 GB, 2048 sector size				
8105	52 pack, CCW WORM 5.2 GB, 2048 sector size				
8004	10 pack, CCW WORM 4.8 GB, 1024 sector size				
8104	52 pack, CCW WORM 4.8 GB, 1024 sector size				
8497	10 pack, CCW WORM 2.6 GB, 1024 sector size				
8597	52 pack, CCW WORM 2.6 GB, 1024 sector size				
8009	10 pack, Permanent WORM 5.2 GB, 2048 sector size				
8109	52 pack, Permanent WORM 5.2 GB, 2048 sector size				
8517	10 pack, Permanent WORM 2.6 GB, 1024 sector size				
8617	52 pack, Permanent WORM 2.6 GB, 1024 sector size				

13.6 IBM 3466 Network Storage Manager

This section covers the IBM 3466 Network Storage Manager storage management solution available from IBM:

Figure 91 shows the IBM 3466 Network Storage Manager.



Figure 91. IBM 3466 Network Storage Manager

Embodying the Seascapes architecture, this integrated network-attached storage solution can provide enterprise-wide data backup for hundreds of workstations, file servers, and desktop machines. The Network Storage Manager automates backup, recovery, storage management, and disaster recovery for all major distributed platforms, including Sun, Windows NT, HP, DEC, Novell, Apple, SGI, and all other ADSTAR Distributed Storage Manager (ADSM) supported platforms.

- Integrates award-winning ADSM storage management software with high-performing SSA disks
- Offers a choice of tape technologies
- Minimizes network traffic through incremental backups and data compression
- Provides an automated, “lights out” network storage management solution for over 30 platforms
- Provides data archiving

- Provides Hierarchical Storage Management
- Provides disaster recovery, planning, and execution
- Enables online incremental backup of DB2, Oracle¹, Sybase¹, Informix¹, Lotus, Microsoft, and SAP/R3¹ databases
- Provides administrative function through a familiar Web browser interface using ADSM WebShell

Table 158 summarizes the key features of the 3466 Network Storage Manager.

Table 158. IBM 3466 Network Storage Manager at a Glance

Model	C00	C10	C20	C30
Max. Tape Drives	NA	4 ²	12 ³	6 ³
Max. Cartridges	NA	30	648	1295
Max Tape Capacity				
Uncompacted	NA	1.0 TB	2.7/4.8 TB ⁴	14.8 TB
Compacted	NA	2.1 TB ⁵	5.4/9.6 TB ^{4, 6}	38.8 TB ^{4, 6}
Disk Storage	36 to 432 GB Serial Storage			
Tape Technology	Refer to Ordering Highlights	Integrated DLT Library	Magstar L18/L32 Library ⁴	MP Magstar 3494 Library
Network Attachment	1-4 of Fast Ethernet, FDDI, ATM, token ring			
Software Pre-req.	IBM Network Storage Manager Software Package V3R1			
Warranty Period	One-year limited warranty			
Physical Specifications				
Model	C00	C10	C20 ⁴	C30
Height	70.78" 1800 mm	40.5" 1029 mm	59.8/70.9" 1518.2/1800 mm	70.9" 1800 mm
Width	28.5" 723.9 mm	39.7" 1009 mm	39.7/29.5" 1009/750 mm	29.5" 750 mm
Depth	41" 1041.4 mm	33.9" 861 mm	33.9/60" 861/1525 mm	60" 1525 mm
Weight	905 lb 410.4 kg	347 lb 157.5 kg	504/832 lb 228.5/337.2 kg ⁷	1233 lb 529.3 kg ⁸

Model	C00	C10	C20	C30
Ordering Highlights The Model C00 does not include an integrated tape library. The Model C00 customer is responsible for selecting a supported tape library and integrating it with the Model C00. This is most easily accomplished through buying these services from IBM or an IBM Business Partner.				

Notes:

1. Using third-party offering
2. Maximum of two tape drives per DLT library
3. Maximum of six tape drives per tape library
4. C20 with Magstar MP L18 Library/C20 with Magstar MP L32 Library
5. 2:1 Compression Ratio
6. 3:1 Compression Ratio
7. With two drives, no cartridges
8. Without cartridges

Chapter 14. AIX and the Bonus Pack on the RS/6000

This chapter discusses the AIX operating system and the AIX Bonus Pack.

14.1 The Base Operating System - AIX

This section will describe the Advanced Interactive Executive (AIX) operating system. A description of hardware and software requirements, release and support dates, new AIX 4.3.0 through 4.3.3 features, and support information will be outlined.

14.1.1 Hardware and Software Requirements

The following section discusses the hardware and software requirements.

14.1.1.1 Hardware Requirements

AIX will operate on IBM Power, POWER2, Power Series 830 and 850 desktop systems, PowerPC systems, POWER3 systems, or RS/6000 SP and SP2 systems with the following exceptions.

- RS/6000 7016 POWERserver, Model 730
- RS/6000 7007 Notebook Workstation Model N40
- POWER Network Dataserver 7051
- RS/6000 7249 Models 851 and 860
- RS/6000 7247 Models 821, 822, and 823

14.1.1.2 New Hardware Support

New hardware support is available for the following systems.

- RS/6000 Enterprise Server Model S80 with up to 24-way processors and 64 GB real memory
- RS/6000 Enterprise Server 7046 Model B50
- RS/6000 Model 7025-F80
- RS/6000 Models 7026-H80 and M80

The Models F80, H80, and M80 require AIX 4.3.3-03 recommended maintenance package (APAR IY09047), which is included on all pre-installed systems and on the 04/2000 Update CD that ships with AIX 4.3 as of 04/2000.

In addition, there is APAR IY09814, that includes additional fixes that are not available before the 4.3.3.0-03 package is shipped. To install to these models from CD, AIX 4.3.3 CD dated 04/2000 (LCD4-0286-05) or later is required. Otherwise, these systems will not boot from older AIX 4.3.3 CDs. These maintenance levels can also be downloaded from the Internet to install using NIM. One link where maintenance is available from is:

<http://techsupport.services.ibm.com/rs6k/fixes.html>

AIX 4.3.2 included hardware support for these models and devices:

- Diagnostics for Models 701X-S7A Enterprise Server S70 Advanced and Model S80, using error log analysis, for the I/O drawer environmental and power warnings that are recorded in the system error log.
- 2 GB Memory cards for RS/6000 7017-S7A Enterprise Server S70 Advanced (which requires a firmware update). AIX supports up to 64 GB of real memory on the S80 was introduced in AIX Version 4.3.3 (the Model S70 is now withdrawn).
- RS/6000 43P 7043 Model 260, a design and analysis workstation or entry commercial server.
- RS/6000 43P 7043 Model 150, an entry-level PCI workstation or entry commercial server.

14.1.1.3 OpenGL and GL 3.2 or graPHIGS Supported Hardware

OpenGL and GL 3.2 for AIX, Version 4.3 or graPHIGS for AIX, Version 4.3 support the following machines:

- RS/6000 7006, all models
- RS/6000 7009, all models
- RS/6000 7011, all 200 series models
- RS/6000 7012, all models
- RS/6000 7030, 3AT, 3BT, 3CT
- RS/6000 7013, all models
- RS/6000 7020, 40P
- RS/6000 7025, Model F40
- RS/6000 7043, all models
- RS/6000 7248 (Series 43P), all models
- RS/6000 Power Series 830 (PHIGS only)
- RS/6000 Power Series 850 (PHIGS only)

The machines listed above must be equipped with one of the following graphics adapters to utilize OpenGL and GL 3.2 for AIX or PHIGS for AIX functions. These adapters below preceded by an asterisk also support GL 3.2. Those preceded by a + support Easy MP (the IBM internally multithreaded implementation of OpenGL) and standard OpenGL.

- POWER Gt3
- POWER Gt3i
- *POWER Gt4e
- *POWER Gt4
- *POWER Gt4i
- *POWER Gt4x
- *POWER Gt4xi
- POWER GXT150
- POWER GXT150L
- POWER GXT155L
- POWER GXT150M
- POWER GXT150P
- POWER GXT250P
- POWER GXT255P
- *POWER GXT500
- *POWER GXT500D
- *POWER GXT500P
- *POWER GXT550P
- *POWER GXT800P
- *POWER GXT800PT
- *POWER GXT800M
- *POWER GXT1000
- POWER GXT2000P
- POWER GXT3000P

Graphics Displays supported by the above listed graphics adapters include:

- IBM 6091 machine type displays
- IBM POWER displays

- IBM color monitors

14.1.1.4 Software Requirements

Table 159 outlines the memory and disk requirements of the various software. The sections that follow discuss these requirements in detail.

Table 159. Memory and Disk Requirements

Software	Memory Requirements	Disk Requirements
AIX 4.3.3	Minimum 32 MB up to 64 GB on the Model S80	315 MB Disk plus at least 64 MB of paging space = 380 MB total (does not include online documentation and additional application software).
OpenGL and GL 3.2	Minimum 32 MB	10-120 MB
graPHIGS for AIX 4.3	Minimum 32 MB	10-120 MB
AIX Fast Connect	Minimum 16 MB plus 1 MB for every connected client	50 MB

Disk Space

Most applications will require additional disk storage space for application data as well as the application program itself. Users should plan ahead for growth of user data and future application needs. As with any UNIX system, performance and capacity of the system will usually benefit from additional memory. The amount of benefit, if any, depends on the applications and system load. Increased memory sizes also imply that a larger amount of disk storage space be reserved for paging space. A typical recommendation is to reserve paging space two to three times the size of real memory. Different rules apply to larger memory systems (approaching 1 GB or more) where paging space may not be as important. Some applications may benefit from tuning the size of the paging space.

Paging Space

A deferred paging space algorithm significantly decreases the size of paging space required for systems with large physical memory. Versions of AIX before 4.3.2 allocated paging space blocks for pages of memory as the pages were accessed. On a large memory machine, where the application set is such that paging is never or rarely required, these paging space blocks were allocated but never needed. AIX Version 4.3.2 and following releases implement deferred paging space allocation in which the paging space blocks

are not allocated until paging is necessary, thus helping reduce the paging space requirements of the system.

Secureway Directory Server

AIX 4.3.1 with APARs IX72439, IX74821, IX75022, and PTF U457544, or AIX 4.3.2 or 4.3.3 are required for Secureway Directory Server.

The directory server administration requires one of the following listed Web servers (or later version) installed on the system:

- Apache 1.3.2, or later
- Lotus Domino Go Server 4.6.2, or later
- Lotus Domino Server 5.0.1, or later
- Netscape FastTrack Server 3.01, or later
- Netscape Enterprise Server 3.5.1, or later
- Microsoft Internet Information Server (IIS) 2.0
- HTTP Server 1.3.3.1

Web Browser

A Web browser must be installed for use with the Web administration features of the server. The browser must be frame-enabled and support AIX Developer Kit, Java Technology Edition, including the IBM implementation for AIX of Sun's Java Version 1.1 AWT events and HTML Version 3.0. The following are the Web browser requirements for this specification:

- Netscape Navigator/Communicator 4.06, or later, for AIX
- A minimum of 64 MB RAM or more is recommended
- Disk Space required as follows:
 - If you already have DB2 installed, approximately 25 MB is needed to install the product.
 - If you do not have DB2 installed, approximately 70 MB is needed to install the product.

14.1.2 AIX 4.3 Releases and Support Dates

Table 160 indicates the release dates of AIX Version 4.3 and their support dates.

Table 160. AIX 4.3 Version Availability Dates

AIX Modification Level	Availability Date	Customer Service Withdrawal Date
4.3.0	October 31, 1997	December 31, 2001
4.3.1	April 24, 1998	December 31, 2001
4.3.2	October 23, 1998	December 31, 2001
4.3.3	September 19, 1999	December 31, 2003

14.1.3 AIX Version 4 Release 3 Features

AIX Version 4 Release 3 is a UNIX-based operating environment designed to handle the needs and requirements of a wide variety of systems and applications. It is designed for use in technical and commercial environments and to be scalable on a wide variety of hardware platforms. AIX Version 4.3 is a cumulation of the functions introduced at all previous releases (unless they have been withdrawn), and those features will not be outlined here - only the new features and functions introduced at AIX Version 4.3.0 through 4.3.3 are discussed.

AIX Version 4.3 support can be characterized by the following categories:

- Processors and application support
- Graphics and GUIs
- Conformance to major industry standards
- International language support
- Operating system functions and utilities system management facilities
- Security and control facilities
- Communication and utilities
- System performance enhancements
- System security
- Network security
- PC interoperability - AIX FAsT Connect
- Java

- Additional printer support

14.1.4 Functions and Utilities - System Management

This section outlines the enhancements and abilities of the following categories:

- Support for uniprocessor and 24-way SMP Scalability
- Memory management
- Local run queues for runnable threads
- Limits on files and threads
- Volume group changes
- Kernel
- AIX Workload Manager
- AIX console logging
- Alternate Disk Installation
- Web-Based System Manager
- Dynamic processor deallocation
- PCI Hot-plug Adapters
- System Backup on CD
- Converged install and NIM scalability
- Semaphores, shared memory regions and message queues
- ODM
- The binder library
- Improved scheduler responsiveness
- Remote reboot
- License Use Manager (LUMs)
- AIX Version 4.3 documentation
- Unified documentation library services
- 32-bit and 64-bit system and application support
- Logical volume manager and disk management and RAID
- Online JFS backup
- Support mirroring for system dump
- AutoFS

- Error messaging
- Fast single instruction patch
- Pthread debug library

14.1.4.1 Support for Uni-processor and 24-Way SMP Scalability

AIX 4.3.3 boasts superior kernel scalability enhancements. These enhancements nearly triple online transaction processing throughput while only requiring a doubling of the number of CPUs (12 to 24), memory (32 GB to 64 GB), and increases in processor speed. This increased throughput is especially meaningful in Enterprise Resource Planning (ERP) and Online Transaction Processing (OLTP) applications. Individual results may vary due to different applications and workload combinations.

14.1.4.2 Memory Management

Memory management has been improved in AIX Version 4.3.3 to accommodate memory of 32 MB to 64 GB to allow higher concurrence with multiple frame lists and multiple page replacement daemons. This reduces contention in the serialization mechanisms and allows processes with lower latencies to service the memory requests.

Replaceable Malloc

AIX now allows the use of alternative memory allocation (malloc) routines using an environment variable. This provides the ability to fully replace the system malloc routine, including the one used within libc itself, without any relinking required. Previously, a malloc replacement would only be used for application calls directly linked with a new malloc but would not be utilized by shared libc routines, possibly resulting in undetectable memory references. This enhancement helps make memory references more visible, enabling the third-party memory access products performing functions, including garbage collection, memory usage analysis, and high speed memory allocation.

14.1.4.3 Local Run Queues for Runnable Threads

AIX Version 4.3.3 organizes the runnable threads into per-CPU local run queues. This simplifies the process of determining which thread to run next and eliminates the costly logic, which was necessary to promote good cache affinity when scheduling from a single global run queue. It also virtually eliminates locking contention in this performance-critical subsystem.

As Workload Manager (WLM) enables greater system loads, some contention is found in non-WLM areas. If you work with large thread loads and use the `svmon` command, you need APAR IY06125. It can be downloaded from:

<http://service.software.ibm.com/support/rs6000>

AIX Version 4.3.3 is a pre-requisite requirement.

With local run queues, the dispatcher's affinity algorithms have been strengthened, resulting in greatly improved throughput on busy SMP systems. User mode threads generate less cache interference as a result of their increased affinity. These reductions in system overhead translate directly into increased application throughput.

Enhancements to AIX Version 4.3.1 improved the scalability of AIX and the solutions implemented upon it.

14.1.4.4 Limits on Files and Threads

The limits on files and threads have been extended to:

- The number of open files per process (ulimit -n) can now be set to a maximum of 32,767. The default limit remains 2000.
- A total of one million open files per system and 32,767 threads per process.
- Maximum size of 2 GB for a shared file or mapped file.
- Restriction of 1016 physical partitions per disk is removed.
- Since AIX 4.2 larger than 2 GB to a maximum file size of 64 GB. The maximum file system size is 1 TB.

14.1.4.5 Volume Group Changes

New volume group limits are a combination of the number of physical partitions per disk and the number of disks in a volume group. Examples of ways to define a volume group are:

- 32 disks with 1016 physical partitions per disk
- 1 disk with 1016*32 physical partitions
- 16 disks with 1016*2 physical partitions per disk

Current volume groups must be converted to take advantage of the new limits. You are not required to convert existing volume groups.

14.1.4.6 Kernel

The basic kernel functions can be described as:

- I/O
 - Device I/O
 - File I/O
 - File systems

- Process Management
 - Load and execution
 - Process scheduling
 - Memory management
 - Interprocess Communications (IPC)

The kernel is further enhanced through kernel extensions.

A kernel extension provides additional kernel services and system calls by supplying an export file when it is link-edited. The kernel provides the set of base kernel services to be used by the kernel extensions.

Since the AIX kernel is a dynamically extendable kernel, it can be expanded by adding device drivers, system calls, kernel services, or private kernel routines.

Enhancements

The kernel trace buffers are no longer required to be in the kernel heap. You can now allocate one or two trace buffers up to 268, 435, or 184 bytes.

The printer queue IDs are expanded from 3 to 6 digits

Kernel Data Addressability

The maximum data space for kernel data and communications buffers was increased in AIX Version 4.3.2. Kernel space for global data structures (kernel heap) was increased from 256 MB to 512 MB to handle much larger system load requirements. This enhancement enables higher throughput for more processes and threads as well as providing ISVs with more space for kernel extensions and device drivers. Communications buffers were expanded from 128 MB to 1 GB to dramatically increase the ability to handle larger communications workloads such as required by Web servers.

Kernel Debugger

The Kernel Debugger, `kdb`, is a tool being added to the AIX system to provide a symbolic debugger for the AIX kernel, kernel extensions, and device drivers. `kdb` is also a command to allow examination of system crash dumps. `kdb` is an alternative to the current kernel debugger and crash command.

14.1.4.7 AIX Workload Manager

The AIX Workload Manager (WLM) introduced with AIX Version 4.3.3 provides a policy-based method for managing system workload and system resources. AIX Workload Manager includes the following capabilities:

- Defines system resource allocations that can be applied towards specific jobs or job classes.
- The operating system allocates CPU and memory resources to jobs or job classes in accordance with the defined resource allocation policies.
- Helps ensure that critical applications are not impacted by less important jobs in the system during peak demand.
- Allows logical job separation on the server.
- Permits applications to remain in memory for more predictable performance.
- Helps provide greater convenience and control by using both resource targets and resource limits.
- Allows policies to be set by the system administrator once, with no further interaction required. The system will automatically apply the specified policies and adjust for changing conditions.
- Permits creation and management of 29 classes of jobs, each with different resource policies and system administrator specified names.
- Allows creation of automatic classification rules to assign processes to classes.
- Permits usage of nine tiers of jobs, with each tier's resource needs being satisfied before resources are provided to jobs in the next tier.
- Provides control options that include minimum and maximum percentage limits, shares, or a combination of both.

14.1.4.8 AIX Console Logging

AIX now treats system console messages as critical system information. Previously, these messages were simply displayed on the current console device. If that screen or window was in use, the messages could be lost. Now, in addition to displaying them on the console, these messages are also logged to a file along with the originating user and the time of day when the message was written. Now, customers can easily retrieve these messages, thus improving their ability to diagnose problems and monitor system status.

14.1.4.9 Improved Serviceability Features

In addition, AIX has numerous new and improved serviceability features. The file system and system dump processor have been enhanced to make it easier to diagnose problems when they occur.

14.1.4.10 Alternate Disk Installation

Alternate Disk Installation was introduced in AIX Version 4.3.0. It allows a customer to install a new release or maintenance level of AIX on an alternate disk without taking the system down for an extended period of time. This is important in 7 x 24 operations and for sites with a large amount of systems.

The AIX command to perform this function is:

```
alt_disk_install
```

It installs an alternate disk with a backup install image or clones the currently running system to an alternate disk.

14.1.4.11 Web-Based System Manager

Web-Based System Manager was originally designed to provide the capability to manage AIX systems locally or remotely, such as from anywhere on the Internet/intranet.

It was introduced in AIX Version 4.3.0 for use with pre-production environments.

With AIX Version 4.3.1, Web-based System Manager was first considered to be fully functional.

With AIX Version 4.3.2, a new application is added to the Web-based System Manager. It allows you to format and certify some hard disks and read-write optical media. You can also run hardware problem determination using an intuitive, object-oriented, and easy-to-use interface. A new feature added to hardware problem determination allows you to specify the number of days to go back for the error log analysis. The range is from one to 60 days. The default value is seven.

In addition, AIX Version 4.3.2 function is extended to provide the capability of launching HTML-based applications, such as Netfinity Manager, from an AIX window using the AIX Web-based System Manager launchpad. The application to be launched can be registered from the command line or with an AIX supplied GUI. By registering Netfinity Manager, this launch capability provides access to all Netfinity Service Web Interfaces. For more information on Netfinity Manager, visit the following Web site:

```
http://www.pc.ibm.com/us/netfinity/system\_management.html
```

With AIX Version 4.3.3, Web-based System Manager is a fully matured and functional system management tool for managing AIX systems. It is designed and implemented to run in application mode on AIX with Java 1.1. It is also

designed to run in applet mode on platforms that support a Java 1.1, AWT 1.1-enabled Web browser. For example, users can manage an RS/6000 system from an AIX client running Hot Java. Users can also manage AIX systems from a PC running Windows 95/98/2000 or Windows NT with Java 1.1, AWT 1.1-compliant browsers, such as Hot Java, appletviewer, and Netscape Communicator 4.02 with Preview Release of JDK 1.1 patch.

The Web-based System Manager features an intuitive, object-oriented, easy-to-use graphical user interface. In addition to a Java-based GUI, it provides TaskGuides to simplify unfamiliar or otherwise complex tasks. The Web-based System Manager provides applications for all major AIX system management tasks. Included are applications for managing:

- Devices
- Print queues
- Communications
- File systems
- Logical volumes
- Backups
- Users and groups
- Processes
- Subsystems
- System environments
- Software
- Network installation
- Cultural environment

Applications for managing software and installation are provided for local and remote systems via the Network Install Manager (NIM).

The Web-based System Manager is NLS-enabled.

Web-based System Manager will automatically show the current resource utilization by class whenever the top level of WLM management is shown.

14.1.4.12 Dynamic Processor Deallocation

Otherwise known as Dynamic CPU Deallocation, dynamic processor deallocation was announced as an additional function to AIX Version 4.3.3 in February 7, 2000.

Combining AIX Version 4.3.3 function and support in hardware and firmware, this capability is designed to provide higher levels of system availability on SMP platforms by continuously monitoring CPU health and dynamically deconfiguring any processor displaying indications of an impending failure. If a threshold of recoverable errors is reached, the workload from the faulty processor is reassigned, that processor is dynamically removed from the active configuration, and operations continue on the remaining processors. Processor failures that occur without warning are not handled by this function.

Dynamic CPU Deallocation is supported on the 375 MHz POWER3 SMP Thin/Wide Nodes and the Model 270, F80, H80, M80, as well as on the Model S80. Firmware level 991216 must be installed on the S80.

The typical flow of events for processor deallocation is as follows:

1. The firmware detects that a recoverable error threshold has been reached by one of the processors.
2. AIX logs the firmware error report in the system error log, and, when executing on a machine supporting processor deallocation, starts the deallocation process.
3. AIX notifies non-kernel processes and threads bound to the last logical CPU.
4. AIX waits for all the bound threads to move away from the last logical CPU. If threads remain bound, AIX eventually times out (after ten minutes) and ends the deallocation.
5. Otherwise, AIX invokes the previously registered High Availability Event Handlers (HAEHs). An HAEH may return an error that will end the deallocation.
6. AIX continues with the deallocation process and ultimately stops the failing processor.

In case of failure at any point of the deallocation, AIX logs the failure with the reason why the deallocation was aborted. The system administrator can look at the error log, take corrective action (when possible), and restart the deallocation. For example, if the deallocation was aborted because at least one application did not unbind its bound threads, the system administrator could stop the applications, restart the deallocation (which should go through this time), and restart the application.

Turning Processor Deallocation On and Off

Dynamic Processor Deallocation can be enabled or disabled by changing the value of the `cpuguard` attribute of the ODM object `sys0`. The possible values for the attribute are `enable` and `disable`.

The default, in the current version of AIX, is that the dynamic processor deallocation is disabled (the attribute `cpuguard` has a value of `disable`). System administrators who want to take advantage of this feature must enable it using either the Web-based System Manager system menus, the SMIT System Environments menu, or the `chdev` command.

Note

If processor deallocation is turned off, AIX still reports the errors in the error log, and you will see the error indicating that AIX was notified of the problem with a CPU (`CPU_FAILURE_PREDICTED`).

Restarting an Aborted Processor Deallocation

Sometimes the processor deallocation fails because, for example, an application did not move its bound threads away from the last logical CPU. Once this problem has been fixed, by either unbinding (when it is safe to do so) or stopping the application, the system administrator can restart the processor deallocation process using the `ha_star` command.

The syntax for this command is:

```
ha_star -C
```

where `-C` is for a CPU predictive failure event.

Processor State Considerations

Physical processors are represented in the ODM database by objects named `procn`, where `n` is the physical processor number (`n` is a decimal number). Like any other device represented in the ODM database, processor objects have a state (`Defined/Available`) and attributes.

The state of a processor object is always `Available` as long as the corresponding processor is present, regardless of whether it is usable by AIX. The state attribute of a processor object (`proc`) indicates if the processor is used by AIX and, if not, the reason. This attribute can have three values:

- `enable` - The processor is used by AIX.
- `disable` - The processor has been dynamically deallocated by AIX.
- `faulty` - The processor was declared defective by the firmware at boot time.

In the case of CPU errors, if a processor for which the firmware reports a predictive failure is successfully deallocated by AIX, its state goes from enable to disable. Independent of AIX, this processor is also flagged as defective in the firmware. Upon reboot, it will not be available to AIX and will have its state set to faulty. But, the ODM proc object is still marked Available. Only if the defective CPU was physically removed from the system board or CPU board (if it were at all possible) would the proc object change to Defined.

The only limitation is that you need a minimum of three processors. The reason you need three is to prevent going from SMP to uni-processor because of locking tests for number of CPUs - if the system went from two processors to one, some drivers and perhaps the kernel could become confused. Deallocation can continue until there are two processors remaining.

14.1.4.13 PCI Hot-Plug Adapters

On September 13, 1999, IBM announced hot-plug support for a broad range of PCI adapters first in the SP Expansion I/O Unit. These adapters include storage and connectivity functions, such as SSA, Ultra SCSI, Ultra2 SCSI, Gigabit Fibre Channel, Ethernet (10/100 Mbps and Gigabit), token ring, and S/390 ESCON attach.

AIX Version 4.3.3 provides interfaces that allow the user to list the hot-plug slots and to add, remove, or replace hot-plug PCI adapters without having to power off or reboot the system.

The hardware that can support the hot-pluggable PCI slots include:

- SP Expansion I/O Unit
- 7025-F80, 7026-H80, 7026-M80

Some sample commands are:

- `lsslot -c pci` (lists slots and their characteristics).
- `drslot` (dynamically reconfigure slots), then run `cfgmgr`.
- `smitty devdrpci` (hot plug manager).
- Web-based System Manager can also be used to manage PCI slots.

If you are running HACMP on a system that allows Hot-Plug removal and replacement of adapters, special care must be taken when replacing network adapters that are part of an active HACMP cluster configuration.

14.1.4.14 System Backup on CD

AIX Version 4.3.3 provides a new capability to generate an AIX system backup (mksysb) onto recordable CDs using third-party writers and third-party software. This function also allows system administrators and Independent Software Vendors to produce customized AIX system installation media on CDs as an alternative to the AIX Network Installation Manager (NIM).

To make this new function easy to use, a TaskGuide is provided to walk the user through the steps needed to create a system backup on CD-ROM. Refer to the Install Guide for more details on tested configurations of third-party hardware and software.

14.1.4.15 Converged Install and NIM Scalability

The Network Installation Manager (NIM) in AIX Version 4.3.2 is enhanced to provide more control for system administrators in handling system installations and customization of NIM groups. This release allows administrators to limit the number of machines in a group that can be installed or customized at a single time. This gives administrators the ability to control network activity while exploiting the benefits of grouping machines. Additionally, this option lets administrators specify a time limit for processing the operations on a group, further enhancing the system administrator's control through NIM.

AIX Version 4.3.3 provides some building blocks to begin to unite the install solutions across the RS/6000 family. Improvements are being put in place in the network installation manager (NIM) by adding scalability features, such as replication of resources and improved NIM master process handling. Base operating system installation (for both stand-alone and NIM installations) has been made more usable with TaskGuides in the Web-based System Manager Tool. Updating stand-alone AIX systems or SP is easier with the introduction of the Software Maintenance TaskGuide, which analyzes and applies updates to your system.

14.1.4.16 Semaphores, Shared Memory Regions, Message Queues

The interprocess communication subsystems are enhanced in AIX V4.3.2 to support significantly more usage per system. The number of semaphores, shared memory regions, and message queues have each been increased to 131,072 items per system.

14.1.4.17 ODM

The Object Data Manager (ODM) customized database, as of AIX 4.3.2, was expanded to enable system configurations up to 1,000 devices.

14.1.4.18 The Binder Library

The binder library (libld.a) supports operations on 32-bit objects from 64-bit programs and operations on 64-bit objects from 32-bit applications.

14.1.4.19 Improved Scheduler Responsiveness

With AIX Version 4.3.2, the scheduler can place more emphasis on the value of the AIX `nice` command, treating the value as a directive instead of a suggestion. With AIX Version 4.3.2, running the process in background or increasing the value of `nice` will result in the process getting significantly less CPU time if it is a heavy CPU user on a heavily loaded system.

14.1.4.20 Remote Reboot

AIX Version 4.3.2 provides a configurable facility to reboot the system using an integrated serial port at interrupt level priority. It can be used to reboot a system remotely when the system is capable of responding to interrupts but is otherwise inaccessible. External security should be provided for the port itself.

14.1.4.21 License Use Manager (LUM)

On November 8th 1999, a new release of the License Use Management product, LUM 4.5.5, was released worldwide. LUM is used to release a certain number of licenses of an applications to stand-alone or networked users.

For more information, go to:

<http://www-4.ibm.com/software/is/lum/support.html>

14.1.4.22 AIX Version 4.3 Documentation

AIX Version 4.3 documentation is designed to be viewed using Web-based browsers. The documentation is shipped as part of the operating system on the following CDs:

- AIX Version 4.3 Base Documentation CD
- AIX Version 4.3 Extended Documentation CD

The AIX Version 4.3 Base Documentation CD contains the following types of documentation:

- User guides
- System administrator guides
- Application programmer guides
- Web-based System Manager extended helps

- All command reference volumes
- Files reference
- Technical reference volumes that application programmers use

The AIX Version 4.3 Extended Documentation CD contains the following types of documentation:

- Adapter guide and reference books
- Technical specifications describing industry standard functions
- Guide and reference books primarily for the system programmers use

The system administrator can mount the Documentation CDs or can copy the contents of the CD onto a system. The Extended Documentation CD included with AIX Version 4.3.1 is now installable as well as mountable.

Most of the AIX Version 4.3 documentation, including the libraries' home pages and navigation articles, is designed to be viewed using forms-capable, HTML Version 3.2-compliant Web browsers. For easiest access to information, the user should choose a browser that supports frames, such as the Netscape Navigator 4.0 browser that ships with the AIX 4.3 Bonus Pack. Search indexes are supplied for HTML-based, single byte language (Latin 1 code set: ISO8859-1 and IBM-850) documentation, making most of the AIX Version 4.3 libraries fully searchable. A user can perform general searches by searching all indexes registered on the system, or the user can control searches by selecting specific indexes to search.

A small portion of the AIX Version 4.3 documentation is shipped in PDF format. This documentation can be searched and viewed using the Adobe Acrobat Reader that is shipped with the AIX 4.3 Bonus Pack. With the enhanced AIX Version 4.3 `man` command support, the reference documentation shipped in HTML format can be accessed either through a browser or through the `man` command.

Much of the documentation that is shipped with the AIX Version 4.3 operating system is available to Internet users through the IBM RS/6000 Resource home page at:

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

AIX Documentation Search

The AIX Version 4.3.2 Documentation Search is extended to add the capability to search specified double byte character set (DBCS) codesets in Japanese, Korean, Simplified Chinese, and Traditional Chinese as in the following table:

- Language CCSID Codeset Locale
- Japanese 932 IBM-932 Ja_JP
- Korea 970 IBM-eucKR ko_KR
- Simplified Chinese 1383 IBM-eucCN zh_CN
- Traditional Chinese 950 big5 Zh_CN

In addition to the search capability, you can also register HTML documents written in the above languages and codesets.

14.1.4.23 Ease-of-Use

Enhanced ease-of-use capability provides SMIT and Web-based System Manager interface support for:

- NIS+
- SecureWay (directory exploitation — users and group information)
- Logical Volumes (RAID 0 + 1)
- AIX Workload Manager
- Quality of Service (Differentiated Quality of Services and RSVP)
- TaskGuides including:
 - Base operating system install
 - Software update
 - System backup on CD-ROM

New Documentation Library GUI offers easier access to online documentation with a single integrated GUI that allows users to read, navigate, and search online HTML documentation.

14.1.4.24 Unified Documentation Library Services

Beginning in AIX Version 4.3.3, the Documentation Library Service is extended to integrate the navigation, reading, and searching of online documents. To use these functions, a new Documentation Library GUI is available. The new documentation Library GUI offers easier access to online documentation with a single integrated GUI that allows users to read, navigate, and search online HTML documentation. The AIX operating system documentation can be accessed through this library service. Additionally, system administrators can register locally written HTML documents into the library so that users can go to a single library GUI to access a wide range of documents. For example, non-AIX documents could include online documentation for customer applications and also company policies and

procedures. The library services can be made available locally, or through use of a Web server, and the documents can be used remotely by AIX or PC computer clients.

14.1.4.25 32-Bit and 64-Bit System and Application Support

The question a lot of customers are asking is when can they move from 32-bit to 64-bit? The move can happen today, with full 64-bit exploitation evolving over time. The move to 64-bit computing does not imply that customer applications must be ported to 64-bit. The performance improvements and business potential of 64-bit computing can begin to be exploited now.

64-bit Systems Support

AIX Version 4.3 supports the new 64-bit Models, such as the Models S80, M80, H80, F80, 270, 170, 222 MHz POWER3 SMP High Node, 375 MHz POWER3-II SMP Thin and Wide Nodes, and 200 MHz POWER3 Thin and Wide Nodes.

These systems can run both 32-bit applications and 64-bit applications. Full binary compatibility is supported for 32-bit applications running on 64-bit systems. In Appendix J, “The RS/6000 64-Bit Solution” on page 755, you will find a detailed document on 64-bit operating systems.

64-bit Application Support

Applications that are likely to perform best from 64-bit computing are:

- Data warehousing
- Data mining
- Decision support
- Internet-based applications
- e-business and e-commerce applications
- Large-scale Web servers
- Multimedia servers
- Numerically-intensive or compute intensive applications

AIX Version 4.3 supports the creation of 64-bit applications on all supported systems and the execution and debugging of 64-bit applications on 64-bit hardware systems. Applications that can make use of large address spaces to house vast amounts of data or map large files, rather than doing explicit I/O operations, can benefit from being ported to a 64-bit environment.

The AIX Version 4.3 assembler, linker, and all XCOFF commands work with XCOFF64, a new 64-bit object format. The 64-bit IBM C for AIX, Version 4.3

compiler supports the creation of 64-bit object files. Most libraries in AIX Version 4.3 support the creation of both 32-bit and 64-bit applications.

32-bit and 64-bit applications have access to equivalent sets of AIX system services, can execute concurrently, and can share access to system resources in the same ways that 32-bit processes currently can.

- The header files in AIX Version 4.3 are modified to provide a common set of header files that support creation of both 32-bit and 64-bit applications.
- The 64-bit programs can make use of multiple threads and large files (since AIX 4.2 larger than 2 GB to a maximum file size of 64 GB. The maximum file system size is 1 TB).
- Files, shared memory, and interprocess communication resources can be shared among 32-bit and 64-bit processes.

AIX Version 4.3 provides support for developing libraries, kernel extensions, and device drivers that will work with 64-bit applications.

Libraries

Non-thread-safe and thread-safe libraries have been combined into one set of libraries, turning thread-safety on by default.

Libraries that are thread-safe include:

- `libc.a`
- `libbsd.a`
- `libm.a`
- `libmsaa.a`
- `libs.a`
- `libdes.a`

These thread-safe libraries enable a convenient programming model for exploiting SMPs and simplify exploitation of threads by applications, middleware, and other Application Programming Interface (API) providers.

The 64-bit C compiler and the AIX Version 4.3 `lint` command provide options to assist in the porting of 32-bit applications to 64-bit applications. These options point out specific areas where porting problems tend to occur.

The 64-bit applications make use of the native 64-bit instructions of 64-bit systems, and the 64-bit C compiler automatically takes advantage of these instructions.

The C base types provided to 64-bit applications by the 64-bit C compiler are compatible with those of most other 64-bit system vendors. The long integer and pointer types are expanded from 32-bit to 64-bit values in accordance with the 64-bit direction commonly referred to as LP64.

All current device drivers continue to work with 32-bit applications, on both 32-bit and 64-bit systems. Device drivers may need to be updated to work with 64-bit applications; all current AIX PCI device drivers have been updated to work with both 32-bit and 64-bit applications on the new 64-bit hardware models.

The ability to execute 64-bit applications can be disabled on 64-bit systems using SMIT and Web-based System Manager.

AIX Version 4.3 adds support for direct I/O, which allows higher bandwidth I/O for performance critical applications.

Note

Since AIX requires the C and C++ run-time libraries to function, they are already packaged in AIX Version 4.3.

C libraries are included in header files, such as libc.a.

C++ V4 RTE is included in AIX 4.3.

VisualAge C++ V5 RTE is not shipped with AIX Version 4.3. However, it is available for download at no charge from IBM.

User Login Performance

User Login Performance in AIX Version 4.3 is significantly improved for systems with a very large number of user accounts. This performance enhancement is achieved with reductions in path length, memory, and CPU usage. On a single-user, low-end system, their performance improvement can be between 50 to 400 percent. This results in improved performance of some of the system commands.

14.1.4.26 Disk Management and RAID 0 + 1

The Logical Volume Manager (LVM) combines RAID 1 (mirror) data availability with RAID 0 (striped) performance by supporting (entirely in software) a striped logical volume with mirrors. This feature further enhances data availability in high-performance striped logical volumes by tolerating disk failures. The remaining disks in the striped mirror copy continue to service striped units contained on these disks. The replacement of a disk where only

the partitions on the new disk are synchronized is provided through the `migratepv` or `replacepv` command.

Super Strict

In addition, all logical volumes now can utilize a new partition allocation policy called Super Strict. This Super Strict policy does not allow partitions from one mirror to share a disk with partitions from a second or third mirror, thus helping to further reduce the probability of data loss with a disk failure.

These new functionalities are not backward compatible; therefore, new volume groups supporting these features cannot be used with previous versions of AIX.

Concurrent Mode and Mirror Resynchronization

Since Version 4.3.0, logical volumes could be created in concurrent mode. The performance of mirror resynchronization in concurrent mode is improved to a level similar to that for non-concurrent mode.

LVM - Enhancements and Backup Tools

LVM enhancements include support for:

- Raw LV online mirror backups is an LVM enhancement that provides a snapshot capability for raw mirrored logical volumes, for example, for Oracle backups. You can use one mirror of a mirrored logical volume to archive the data on the raw logical volume without splitting the mirror copies from each other (only the logical partitions that have changed during the system backup need to be resynced).
- LVM online mirror backup enables improved systems availability to end users by supporting online backup while still providing access to data. In other words, LVM online backup support allows you to backup the LVM without stopping it. Online mirror backup support was introduced in AIX Version 4.3.1. However, this support was limited to a single node, and does not allow file system access to the backup mirror. Also, access to the backup copy was only through special purpose interfaces. In AIX Version 4.3.2, these capabilities are enhanced to provide for concurrent mode support (when using HACMP) and special device file support for the designated online backup copy. This allows for both file system access and access using standard commands and programming interfaces.
- More than 1016 physical partitions per physical volume. This provides support for most volume groups violating the 1016 physical partitions per physical volume limit.
- Physical partition sizes of 512 MB and 1 GB

- LVM scalability in AIX Version 4.3.2 is enhanced with a new volume group format, which increases the maximum number of disks that can be included in a volume group from 32 to 128. The maximum number of logical volumes in this new volume group format is increased from 256 to 512. Volume groups from previous versions of AIX are still supported and a migration path is provided to convert old volume groups to the new volume group format provided there are sufficient free partitions on each of the physical volumes in the volume group.

14.1.4.27 Online File System Backup

The file system is enhanced in AIX Version 4.3.3 to support file system online backup. The capability allows a mirrored copy of a file system to be used for backup purposes. A mirror copy of the file system is split off, mounted, read-only, and available for backup. This enables end users to back up a consistent copy of the file system while another copy is still mounted and in use. After the backup is complete, the user can reintegrate the backup mirror copy and resynchronize it with the other mirror copies.

14.1.4.28 Support Mirroring for System Dump

The restriction that a dump device cannot be a mirrored logical volume has been removed by forcing the dump to be written to and read from the primary mirror of a mirrored logical volume. Previously, if a customer wished to mirror their root volume group, they had to specifically avoid mirroring their dump logical volume. Now, customers can use mirrored paging devices as dump devices. Dump data itself is not mirrored, but the AIX dump management utilities have been enhanced to properly obtain the dump data, regardless of the mirroring state.

14.1.4.29 AutoFS

AIX Version 4.3.1 introduces a port of the Sun ONC+ implementation of automatic mounting, known as AutoFS. This function replaces the current automount facility.

AutoFS consists of three components:

- Automount command
- AutoFS file system kernel extension
- Automount daemon

Because the `automount` command and the automount daemon are now separate entities, the `automount` command can act as an administration tool. This means that the system administrator can update the map files and then inform AutoFS of the changes without stopping and restarting the automount

daemon. This new implementation is not limited to just NFS. You can, for example, automatically mount CD file systems, CacheFS, and more.

14.1.4.30 Error Messaging

As of AIX Version 4.3.2, it is possible to use the standard AIX messaging facility for messages in error log templates. Such templates can specify a message with a message catalog, set, number and default text, or they can still use the existing error logging codepoints. The maximum message length is now 1,000 bytes when standard messages, not codepoints, are used. An additional benefit is that error messages can now take advantage of National Language Support (NLS) to provide a superior customer interface.

This enhancement makes it optional to use separate error logging codepoint files conforming to SNA Alert architecture rules and its message length limits. All error log templates created conforming to SNA Alert architecture rules are still valid and usable. There are no restrictions on mixing the use of existing, shorter messages (40 to 128 bytes) with the use of the new 1,000 byte NLS-enabled messages.

14.1.4.31 Fast Single Instruction Patch

The new Fast Trap signal generates a trap (branch) that will bypass the debugger, permitting users to insert patch or instrumentation code with a single instruction replacement. This results in roughly an order of magnitude improved performance over the current mechanism of using the current trap signal through elimination of notification to the debugger, for example, no context switch required.

14.1.4.32 Pthread Debug Library

AIX Version 4.3.3 facilitates third-party development of user-level, thread-capable debuggers by providing a pthread debug library. This library provides the APIs for the development of debuggers or applications, which require information about both user level threads (pthreads) and kernel threads. Furthermore, the pthread debug library provides the building blocks needed in the development of an application designed to gather information about mutex (lock) and condition variable usage over the lifetime of the application.

14.1.5 Graphics and GUIs

This section will cover the following topics:

- Common Desktop Environment
- X11R5 and X11R6 Windowing System

- Motif 2.1
- OpenGL and GL 3.2 for AIX 4.3
- graPHIGS for AIX 4.3
- Graphics Enhancements

14.1.5.1 GUI based on the Common Desktop Environment (CDE)

The Common Desktop Environment (CDE) Version 2.1 is a graphical user interface for UNIX in its variants (AIX, Digital UNIX, HP/UX, Solaris, UnixWare, and so on). CDE brings unparalleled ease of use to UNIX and is being adopted as a standard operating environment by many companies in the UNIX workstation market. CDE has been an integral part of AIX since Version 4.1.

14.1.5.2 X11R5 and X11R6 Windowing System

The AIX windows environment at AIX 4.3.0 was X11 Release 6.1. The new X Window System Version 11 Release 6 support provided new functions, such as Inter-Client Exchange (ICE), thread-aware libraries, 32-bit and 64-bit libraries, the X Input Method (XIM), the Record Extension, and the X Keyboard Extension (XKB). Font enhancements included a True Type Rasterizer for the Fontserver and enhancements to the X Logical Font Descriptions (XLFD) for matrix-transformed fonts. The AIX proprietary XAsyncInput API was not ported to X11R6.1. The X11R3, X11R4 and X11R5 libraries were provided for backwards compatibility for 32-bit applications.

At AIX V4.3.3, X11R6.3 (also known as Broadway release from the X Consortium) was included to allow a Web server to run X Windows System client applications on a remote host. With this function, system administrators can create simple Web-based menus of applications to run on the server, and users can point and click to run them. For example, sites that have many X Terminals and Network Computers can use this function to provide a single Web-based point of contact for running their enterprise applications.

X11 Release 6.3 includes four new functions:

- Remote/Desktop Agents

Two types of agents are used to facilitate remote execution of X-client applications. They are as follows.

- Netscape plug-ins
- Stand-alone helper applications for the rx mime type to be used with older Netscape versions and other Web browsers
- Security Extension

Introduces the concept of trusted and untrusted applications in the X Windows System. Untrusted applications are not allowed to do unwanted things, such as intercepting keystrokes.

- Application Group Extension

Allows an application to perform some of the functions that a Window Manager normally does, such as putting resize handles on top level windows. This allows an X application to be embedded in a browser.

- Low Bandwidth X (LBX)

Employs techniques to minimize bandwidth, including protocol compressing, re-encoding, short circuiting, and caching to reduce network traffic on slow connections. LBX Proxy allows applications to use LBX without recompiling by acting as a pseudo-server. This means that it is located between the applications and the X Server and passes information between the applications and the X Server.

Virtual Frame Buffer Rendering

Virtual Frame Buffer (VFB) rendering VFB allows you to write Web-based 3-D graphics applications for an RS/6000 server without the need for a 3-D graphics adapter. Users can take advantage of multiple processors for near linear scalability because each client can render into their own frame buffer without interaction with the X server. This is because each client can render into its own frame buffer without interaction with the X Server. The system administrator of the host system must specify `-vfb` on the command line for the X Server or in the `.xserverrc` defaults file for VFB to be used.

Virtual Frame Buffer (VFB) rendering in AIX V4.3.2 is supported on RS/6000 SP systems as well as all other currently supported RS/6000 systems.

End users viewing data at client systems receive the benefit of VFB with no changes to software or invocation methods to applications. Applications must be enhanced to exploit VFB technology.

For more information about VFB, visit the Developer Connection Web site at:

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

14.1.5.3 Motif 2.1

Motif 2.1 was the latest distribution from The Open Group CDE/Motif Project. It included new features, such as Thread-Safe, 32-bit and 64-bit libraries, numerous bug fixes, new widgets, and Complex Text Layout (CTL) support. Motif 1.2 libraries are provided for backwards compatibility for 32-bit applications.

14.1.5.4 OpenGL and GL 3.2 for AIX 4.3

GLX support of Version 1.3, the latest level approved by the OpenGL Architectural Review Board (ARB), is available on the GXT2000 and GXT3000 adapters. New functions include rendering access to off-screen adapter memory (pbuffers) (only available on the GXT2000) and enhanced X visual and resource selection and management.

OpenGL performance enhancements include improved positional lighting and the following new extensions:

- Vertex Array List — An IBM extension that allows a list of vertex attribute arrays to be provided with a single OpenGL function call. This helps reduce performance overhead of multiple function calls.
- Multi-Mode Vertex Arrays — An IBM extension that allows multiple vertex arrays for different primitive types to be provided using a single OpenGL function call. Like the Vertex Array List, this enhancement helps reduce performance overhead of multiple function calls.
- Clip Volume Hint — An extension that allows applications to notify the IBM OpenGL implementation whenever view frustum clipping is not required. This is designed to help reduce rendering computations.

OpenGL Applications Thread Enablement

OpenGL Applications Thread Enablement allows the development of multithreaded OpenGL applications to use direct rendering context where any thread can call OpenGL functions. This is a performance improvement for multithreaded applications since direct rendering contexts have higher performance than indirect contexts.

64-bit Indirect Rendering OpenGL

Users can perform early application development of 64-bit OpenGL clients and remote display of 64-bit OpenGL clients on 32-bit graphics workstations. This 64-bit support is provided for OpenGL libraries and development tools. In AIX Version 4.3.1 64-bit OpenGL supports indirect rendering contexts only. Indirect rendering specifies that OpenGL rendering for the context is performed through the X server.

OpenGL 1.1 and 1.2

OpenGL 1.1 and 1.2 were enhanced in AIX Version 4.3.2. Users of uniprocessor systems will note faster drawing of primitives under most conditions. All users should see improvements in:

- Throughput and cache utilization.
- Latency when lighting is enabled.

- Overall performance on any primitives using the new MultiDrawArray extension.

This release of OpenGL within AIX Version 4.3.2 introduces support for the RS/6000 POWER GXT3000P PCI Graphics Accelerator.

OpenGL Version 1.2 within AIX Version 4.3.2 is enhanced with these new features, which include 3-dimensional textures, additional texture-mapping control, new pixel formatting capability, and support for the Vertex Array Draw Element Range function. Because OpenGL 1.2 is a superset of OpenGL 1.1, all programs written for OpenGL 1.1 run on OpenGL 1.2 without modification, recompiling, or relinking.

OpenGL 1.2 support is available in this release only on the GXT3000P. Users of other graphics adapters are limited to functions contained in OpenGL 1.1 and OpenGL 1.1 extensions. The optional Imaging Extension subset is not supported by the OpenGL 1.2 implementation at the time of AIX Version 4.3.2's announcement

Three new extensions to OpenGL are available in AIX Version 4.3.2. Brief descriptions are as follows:

- MultiDrawArray Extension - Enables users to group together multiple primitives and send them to the adapter with one call. This is supported on all OpenGL-capable adapters, except the GXT1000.
- Texture Mirrored Repeat Extension- Gives users the capability of specifying texture maps without discontinuities at the edges. This is supported only on the GXT3000P.
- Color Blend Extension - Gives users more options in creating blended or translucent colors without having to use an alpha buffer. This is supported only on the GXT3000P.

The ZAPdb Interactive OpenGL Debugger is enhanced to support all new features provided in OpenGL Version 1.2.

Software improvements were made to OpenGL and graPHIGS for AIX Version 4.3.1, resulting in improved performance for certain application scenarios and improved memory utilization by the graphics libraries. Significant performance gains were seen in some SDRC I-DEAS Master Series scenarios. Performance of graPHIGS surfaces, which include CATIA shaded mode, improved by up to 10 percent on the GXT800P accelerator, and in some cases, even more.

14.1.5.5 graPHIGS for AIX 4.3

graPHIGS has been improved with the following performance enhancements:

- Improved CPU utilization during swaps.

The default swapping method of the GXT2000P and GXT3000P graphics accelerators has been changed to minimize the idle time in the graPHIGS API software.

The previous implementation required the device-specific software to be put to sleep while waiting for the vertical blank interrupt from the adapter indicating that the swap has completed. The updated software allows the application to begin working on the next frame before the vertical blank interrupt occurs.

- Improved clipping performance.

Models that are rendered to a view the same size as the graPHIGS workstation window can now use hardware clipping capabilities on the GXT2000P and GXT3000P graphics accelerators to improve clipping performance. Models that extend slightly beyond the window boundaries (up to 20 percent) will no longer be clipped using software. This has shown to significantly improve performance on some models.

- Improved large program support.

In AIX versions prior to AIX 4.3.3, the maximum data addressability for large graPHIGS programs was 1.25 GB (five segments). With changes in AIX 4.3.3, large graPHIGS programs can now access up to 2 GB (eight segments) of contiguous memory.

- Full input and output support of the Japanese IBM 943 encoding.
- As of AIX 4.3.2, graPHIGS is enhanced to improve the performance of polygon rendering under most conditions.
- In AIX V4.3.2, graPHIGS throughput to the graphics adapter is also enhanced, often resulting in performance improvements for the rendering of all primitives.

14.1.6 Standards Overview

This section will list the industry standards that AIX Version 4.3 is based on.

- Base standards
- Binary compatibility
- ISO Standard ISO8859-15
- Tivoli ready

- EuroReady
- Security certification
- VPN certification
- International language support

14.1.6.1 Base Standards

AIX Version 4.3.3 complies with the following standards:

- X/OPEN, including the following standards.
 - UNIX98 Profile Brand
 - XPG4 Network File Systems (NFS)
 - XPG5 Transport Service (XTI) V2
 - XPG5 Sockets V2
- IEEE POSIX, including the following standards:
 - Support for IEEE POSIX 1003.1-1996 (1003.1c) threads option, including M:N thread support

14.1.6.2 Binary Compatibility

AIX Version 4.3.3 continues application Binary Compatibility with previous releases of AIX Version 4.

When the AIX Version 4 binary compatibility rules are followed, applications that are developed and compiled using either AIX Version 4.1 or 4.2 will execute properly on AIX Version 4.3.

For more information, read *AIX Version 4 Binary Compatibility* at:

<http://www.ibm.com/servers/aix/products/aixos/compatibility>

14.1.6.3 ISO Standard ISO8859-15

Support for the ISO standard ISO8859-15 is added to AIX Version 4.3.3. This support provides the ability to process data in the ISO8859-15 codeset, which replaces eight characters from ISO8859-1 with eight new characters, including one Euro symbol, three characters in support of French, and four characters in support of Finnish. All languages that currently use ISO8859-1 will now have the option to use ISO8859-15 instead.

14.1.6.4 Tivoli Ready

In support of developing systems management standards, an optional preload of Tivoli Management Agent can make AIX Tivoli ready. The agent is also included in the AIX Version 4.3 Bonus Pack.

14.1.6.5 EuroReady

IBM considers an IBM product to be EuroReady if the product, when used in accordance with its associated documentation, is capable of correctly processing monetary data in the Euro denomination and of respecting the Euro currency formatting conventions (including the Euro sign). This assumes that all other products (that is, hardware, software, and firmware) that are used with this product are also EuroReady. IBM hardware products that are EuroReady may or may not have an engraved Euro sign key on their keyboards. Currently, EuroReady status applies primarily to IBM products specific to the EMU countries. Products that are not specific to these countries are deemed to be not ready for Euro unless otherwise stated in the product's country-specific specifications.

For more information on the implications of the Euro, visit the IBM Euro Web site at:

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

IBM considers a solution to be EuroReady when the solution providers have:

- Analyzed the Euro requirements, including the need to comply with relevant EC rules
- Built in appropriate function Solution providers can clearly demonstrate this by:
 - Detailing Euro-related requirements
 - Describing how these will be implemented in the solution
 - Declaring when the implementation will be generally available

The Euro capabilities of a EuroReady solution will be clearly identified as features of the Solution specification as described in the contracts document.

14.1.6.6 Security Certifications

AIX Version 4.3 in May 1998, AIX was certified at the high E3/F-C2 level. E3/F-C2 is an international standard, the European Information Technology Security Evaluation Criteria (ITSEC), certified by Bundesamt fuer Sicherheit Informationstechnik (BSI). BSI is a German certification authority.

AIX Version 4.3 was B1 certified before March 1999.

14.1.6.7 VPN Certification

AIX Version 4.3 is the first server operating system to receive ICSA VPN certification. AIX runs on the RS/6000 family of servers and workstations. The certification concerns AIX Version 4.3.0 plus a package of PTFs.

Certification involved rigorous testing to verify that AIX's VPN capabilities meet tough ICSA security standards. The AIX operating system is an element of IBM's eNetwork Virtual Private Networks solutions that help provide safe, flexible, and low-cost end-to-end-encrypted connections across the Internet, and other public networks, that otherwise would be vulnerable to hackers.

14.1.6.8 International Language Support

The basis for international language support is a product provides Unicode, Code Page 943, Bidirectional, Single-Byte Character Set, Multi-Byte Character Set, or UTF-8 Support.

Unicode Support

AIX provides an equivalent set (to the existing AIX locales) of Unicode locales based on UTF-8 (File System Safe Transformation Format for Unicode) as the file code in support of world-wide language requirements. Unicode enables global solutions to code set problems without the limitations of existing national single-byte code sets (SBCS). The SBCS structure, in some cases, is not extensible and does not support the addition of new graphics, such as the Euro currency symbol.

AIX Version 4.3.3 provides IBM Internationalized Classes for Unicode (ICU), a series of programming libraries that will allow application developers to develop C or C++ applications that handle all of the various languages contained within the Unicode standard in a consistent fashion. Furthermore, it enables application developers to write Unicode-enabled applications that will be portable across all IBM operating systems.

Code Page 943

Code Page 943 supports interoperability with Microsoft Windows clients in Japan.

Bidirectional (BIDI)

The following BiDi character sets are supported:

- Arabic
- Hebrew

Single-Byte Character Set (SBCS)

Albanian, Belgian Dutch, Belgian French, Brazilian Portuguese, Bulgarian, Byelorussian, Canadian French, Catalan, Croatian, Czech, Danish, Dutch (Netherlands), Estonian, Finnish, French, German, Greek, Hungarian, Icelandic, Italian, Latvian, Lithuanian, Macedonian, Norwegian, Polish, Portuguese, Romanian, Russian, Serbian Cyrillic, Serbian Latin, Slovak,

Slovene, Spanish, Swedish, Swiss French, Swiss German, Turkish, Ukrainian, U.K. English, U.S. English.

Multi-Byte Character Set (MBCS) support

Multi-byte character set support includes the following:

- Traditional Chinese
- Simplified Chinese
- Japanese (kanji)
- Korean
- Thai
- Vietnamese

UTF-8

UTF-8 character set support includes the following:

- Estonian
- Latvian
- Lithuanian
- Simplified Chinese

Taiwanese Input Method Enhancements

The Taiwanese Input Method has been significantly enhanced for traditional and simplified Chinese Unicode locales. These enhancements have been made to the following input methods for greater similarity to the local Chinese/Taiwanese character inputting methods:

- Intelligent ABC (based on phonetic representation of Chinese characters)
- Biao Xing Ma (in which a Chinese character is divided into several components according to its writing orders)
- Zheng Ma (based on the grapheme representation of a Chinese word)
- Wu Bi (Five Strike — classifies a Chinese character into three categories: Stroke, radical, and single-character)
- Pin Yin (based on the phonetic representation of Chinese characters)

Furthermore, users will have the ability to switch from one input method to another or between the Chinese input method and the English input method.

Korean Input Method Enhancements

In AIX Version 4.3.3, Korean Input Method expands to support all Korean characters defined in KSC 5700. As a result, users can input not only 11,172

HANGEUL characters and 7,744 HANJA characters, but also JAMO (Korean consonants and vowels). The Korean Input Method features the following characteristics:

- Compound consonants and compound vowels can be input as one character.
- Half-width and full-width character input supports ASCII characters in both single-byte and multibyte modes.
- Special characters can be input by conversion from JAMO or code input function.
- An over-the-spot pre-editing drawing area allows intermediate characters in the reverse video area that temporarily covers the text line. The complete character is sent by pressing the conversion key.

Japan Kit Fixed Fonts and Printers Support Integration

AIX Version 4.3.3 integrates various Japanese fixed fonts and Japanese printer support previously provided by separate LPP, Japan Kit V2 (5607-E30), which is available for Japan only. AIX now supports six Japanese font sizes from 6 points to 23 points. Italic and Bold deform fonts are also available for larger font sizes (12, 17, and 23 points).

The Japan Kit V2.1.1 continues to provide the above support for AIX Version 4.2.1. Other components of Japan Kit V2.1.1, for example, Wnn6, advanced Japanese Input Method, and Japanese True Type fonts, are supported by Japan Kit V2 only. Price quotation (PRPQ) will be acceptable upon request to deliver Japan Kit V2.1.1 for countries other than Japan.

New Locales

AIX Version 4.3.3 expands the number of locales supported by four. The newly supported locales are:

- Italian — Switzerland (it_CH or IT_CH)
- English — Australia (en_AU or EN_AU)
- English — Belgium (en_BE or EN_BE)
- English — South Africa (en_ZA or EN_ZA)

These locales support new keyboard maps, such as Low Function Terminal and language conventions, including collation, case conversion, character classification, message catalogs, date-and-time representation, monetary system, and numeric representation. The added locales support will be in both UCS2.0 and ISO8859-15 encodings. Direct translation will not be provided for these languages; however, users and applications can utilize the

existing customary directories to store translated messages, which could be displayed by the new locales.

14.1.7 Communications

This section will outline the new developments in the following areas:

- NIS+
- Sendmail functional update
- LDAP
- Transmission Control Protocol/Internet Protocol (TCP/IP) v6 and v4
- BNU
- Other supported utilities
- TCP/IP, including:
 - Internet Protocol IPv6
 - IPv6 support standards
 - TCP/IP resolver enhancements
 - TCP checksum offload for Gigabit Ethernet
 - TCP/IP socks library
 - IP Security on IPv4 and IPv6
 - TCP/IP gratuitous Address Resolution Protocol (ARP)
- Quality of service
- Network file cache
- IP addresses
- EtherChannel
- DHCP functional update
- NFS3
- WebFS
- PC services - CIFS services
- MPOA on ATM enhancements

14.1.7.1 NIS+

AIX Version 4.3.3 includes a port of the Sun Solaris Version 2.5 NIS+ implementation. This function is provided in addition to the current NIS support, which remains unchanged. This new naming service provides

enhanced capabilities for security management in a distributed system environment. Security can be managed for a set of systems using a single management point. NIS+ was designed to meet the demanding requirements of networks, which typically range from 100 to 10,000 multivendor clients supported by 10 to 100 specialized servers located at various sites. NIS+ enables system administrators to store information about client addresses, security information, mail information, network interfaces, and network services in central locations where all clients on a network can access it. Information is incrementally updated and propagated immediately, allowing information to be changed rapidly. The NIS+ namespace was designed with hierarchical domains to accommodate more distributed networks requiring scalability and decentralized administration. The NIS+ implementation is optimized to support up to 10 replicas per domain. Such a domain may typically have 10,000 table entries. Scalability beyond 1000 NIS+ clients is best achieved by dividing NIS+ name spaces into different domains to create a hierarchy.

Differences between NIS and NIS+ are outlined in Table 161.

Table 161. Differences between NIS and NIS+

NIS	NIS+
Flat Domains: no hierarchy Slave servers on each subnet	Hierarchical layout: Data stored in different levels in the name spaces Support up to 10 replicas per domain. No special requirements for broadcast subnets A domain may typically have 10,000 table entries
Data stored in two column maps tables	Data stored in multicolumn tables
No authentication	Data Encryption Standards (DES) authentication Fine grain access control Access controlled by authentication and authorization
Updates delayed for batch propagation	Incremental updates propagated immediately

14.1.7.2 Sendmail Functional Upgrade

Sendmail has been upgraded to Version 8.9.3, which features anti-spamming among other capabilities. AIX will include the necessary files to generate custom configuration files for the anti-spamming feature. The default `/etc/sendmail.cf`, does not include anti-spamming configuration; however, the

/usr/samples directory contains custom anti-spamming configuration files to illustrate how to configure the anti-spamming feature.

AIX 4.3.1 included Sendmail 8.8.8 that had included LDAP Directory Services and delivery status notification in which it notified the sender when mail was delivered successfully.

14.1.7.3 LDAP

IBM DSSeries LDAP Directory, Version 1.1, that was included with AIX Version 4.3.0, was an initial technology version of the IBM DSSeries LDAP Directory, Version 1.1. This version, and the eNetwork LDAP, Version 1.1.1 in AIX 4.3.1, provided as introductory offerings to enable customers to develop and test Lightweight Directory Access Protocol (LDAP) Directory applications in a pre-production environment. This version is fully supported by IBM through IBMs normal service and support channels. Customer feedback through these channels is welcome.

In AIX 4.3.2, Version 2.1 of the LDAP Directory provides many new enhancements over Version 1.1.1 shipped with AIX Version 4.3.1. The major enhancements include:

- DB2 Version 5.0 as the directory data storage facility
- Alias support
- Improved search and ACL support
- Support for popular Web servers
- Significant scalability
- Improvements to replication and performance

The LDAP Directory now supports up to fifteen million entries with peak sub-second response time for searches. Performance of the LDAP Directory is improved with statement caching and optimization. Multithreading improvements allow LDAP clients to perform true multithreaded connections and make concurrent operations with the DB2 server. The DB2 program provided with AIX may only be used by the eNetwork LDAP Directory function.

The LDAP Directory is based on LDAP Version 2 (RFC 1777, 1778, 1779) plus IETF LDAP Version 3 extensions.

The LDAP Version 2 support is enhanced to include support for aliases.

ACL support provides role-based authorization, assigns multiple users ownership of an entry, and removes the requirement to set ACLs on every node to give users access to their own information.

The Web servers Apache, Lotus Domino Go, Netscape FastTrack, and Netscape Enterprise are supported for LDAP administration.

Client access to the LDAP directory is supported using LDAP or HTTP protocols. AIX client applications can be developed using the enhanced elements provided for supporting LDAP Version 3 protocols and APIs. Also included is the Java Naming and Directory Interface (JNDI) client API that provides Java applications with access to LDAP directories. Both clients support access to LDAP directories using LDAP Version 2 or Version 3.

LDAP client applications can be built for Windows NT, Windows 95, Solaris, and HP-UX using the IBM LDAP Client Pack, which can be ordered separately (PRPQ 5799-GAN).

Also shipped with the LDAP Directory is Directory Sample 1, a client application that creates a directory for testing LDAP. Directory Sample 1 is provided without support.

Beginning with AIX 4.3.3, the administrator can configure the user/group security information to reside in the IBM SecureWay Directory (the IBM brand name for AIX's LDAP implementation). Any number of AIX systems can be configured to retrieve the user/group information from the LDAP, thus sharing the common view and central administration of the user information for that collection of AIX systems.

Interoperability and Compatibility

The eNetwork LDAP Directory replication interoperates with the OS/390 LDAP Server.

The LDAP Directory provides a native, scalable Lightweight Directory Access Protocol (LDAP) directory that can be used by applications and services requiring LDAP standards-based directory services. The directory provided with AIX 4.3.0 is based on the IETF LDAP Version 2 (RFC 1777) along with some extensions for IETF LDAP Version 3. The LDAP directory utilizes IBM DB2 Version 4.1 as the directory data storage facility. DB2 provides a robust, scalable, industry-tested repository for data storage. The DB2 Version 4.1 Single User component is shipped with the LDAP Directory and is installed for use as the directory data storage facility when the directory is installed. The user does not have access directly to DB2.

The LDAP Directory provides Secure Sockets Layer (SSL) Version 3 support, both for the directory server and client. The initial offering is U.S. English only. SSL provides encryption of data and authentication using X.509v3 public-key certificates. The directory can be configured to run with or without SSL support. The directory also supports LDAP referrals, allowing directory operations to be redirected to another LDAP directory server. Replication of the LDAP Directory is supported, which allows additional copies of the directory to be available for directory read operations, thus increasing performance and reliability of access to the directory information. Administration and configuration for the directory is supported by a Web browser-based graphical user interface. The administration and configuration functions provided by the Web-based interface allow the administrator to perform initial setup of the directory, change configuration options, and manage the daily operations of the directory. User access control is provided for the information stored in the directory and can be defined by the administrator. There is also a Web browser interface provided for the user to add, modify, or search for information stored in the directory.

Client access to the directory is supported using LDAP or HTTP protocols. LDAP client applications can be built for various operating system platforms using the IBM LDAP Clientpack that can be purchased separately. An AIX LDAP client is shipped with the LDAP Directory. The AIX LDAP client consists of the LDAP shared library, C header files, sample programs, and documentation for writing programs that use the LDAP application programming interfaces

AIX LDAP client applications can be developed to provide secure access to an LDAP directory server using Secure Socket Layer (SSL) technology. Implementations of SSL, which employ RSA software, allow authenticated access between an LDAP client application and an LDAP server and provide data confidentiality over those connections.

LDAP Standards

- RFC 1777 Lightweight Directory Access Protocol
- RFC 1778 String Representation of Standard Attribute Syntaxes
- RFC 1779 String Representation of Distinguished Names
- RFC 1823 LDAP Application Program Interface
- RFC 1960 A String Representation of LDAP Search Filters

LDAP Software Requirements

A Netscape FastTrack Version 2.0.1 Web server must be co-located on the DSSeries LDAP Directory server for Web-based administration using HTTP

(HTML 3.0/3.2 are supported). Netscape FastTrack Version 2.0.1 is available in the AIX 4.3.0 Bonus Pack.

14.1.7.4 Basic Network Utilities (BNU)

The transfer of files among systems is the most common application of the Basic Networking Utilities (BNU). BNU uses four commands, `uucp`, `uuseed`, `uuto`, and `uupick` to exchange files between local and remote systems.

The `uucp` command is the primary BNU-data transfer utility. The `uuseed` command is the Berkeley Software Distribution (BSD) transfer command incorporated into BNU. The `uuto` and `uupick` commands are specialized to send and receive commands working with the `uucp` command.

The BNU commands, `uuencode` and `uudecode`, assist file transferring. These commands encode and decode binary files transmitted through the BNU mail facility. AIX Version 4.3 continues to provide this support.

14.1.7.5 Other Utilities

The following utilities are supported in AIX Version 4.3:

- Serial Line Internet Protocol (SLIP)
- Point to Point Protocol (PPP)
- Common Data Link Interface (CDLI)
- Data Link Protocol Interface (DLPI)
- Simple Network Management Protocol (SNMP) Agent
- ATM LAN emulation
- IPX/SPX protocols
- Network File Systems (NFS)
- Network Computing System (NCS)
- CacheFS in ONC+ - Adds support from the ONC+ suite. CacheFS is used to cache file data from an NFS mount point in the files in local JFS file system

14.1.7.6 TCP/IP

This section outlines TCP/IP V6 and V4 and their interoperability and security enhancements in AIX Version 4.3.

Internet Protocol Version 6 (IPv6)

IPv6 is the next generation of Internet protocol to address IP address limitation. IPv6 offers robust fail-safe routing, IP address autoconfiguration, and IP security and integrity.

IPv6 implementation of AIX Version 4.3 is a migration platform, where protecting a customer's investment of existing applications is just as important as the introduction of the next generation of IP. IPv4 and IPv6 can be used concurrently on AIX systems. Customer applications written with current IPv4 interfaces will work as is (binary-compatible); however, some of the kernel extension modules may require recompiling to run on AIX Version 4.3. Most of the AIX-provided Internet applications, such as name resolver, configuration, mail, r-cmds (rlogin, rsh), telnet, ftp, and tftp, are upgraded to support IPv6.

AIX Version 4.3 IPv6 provides the required infrastructure for serving all the Internet-ready boxes and devices to participate in the ever-growing Internet. AIX IPv6 implementation supports host requirements. IPv6 IP forwarding is not supported at this time. Major functions provided in IPv6 are described as follows:

- IP Security
IP-layer security services ensure packet authentication, integrity, access control, and confidentiality. Secure IP tunnel between two systems to perform key exchange, message encryption, and message authentication is also provided.
- Neighbor Discovery/Autoconfiguration
With deployment of participating routers, the Neighbor Discovery/Autoconfiguration provides alternative reachable routes in the event of route failure. Combined with address autoconfiguration (based on link address), this addresses the manageability of the IP address and route availability.
- Multihoming
Multiple IP addresses for single host is supported.
- Tunnel support
Two tunnel modes are supported: IPv6 over IPv4 and automatic tunneling to IPv4-compatible addresses.
- Name server enhancement

This enhancement implements BIND Version 8.1.1, including support of IPv6 domain and name API, coexistence with AIX BIND V4 with security update, and support of the AAAA record.

- New application socket programming interfaces to support IPv6 addresses
- Support of IPv6 over Ethernet and token ring

In AIX Version 4.3.2, the BIND package is upgraded to Version 8.1.2 with the following enhancements:

- Improved dynamic updating to the “host name to IP address” database as a result of the addition of Incremental Zone transfer (RFC1995) features.
- Enhancements to user configuration options including notify (RFC 1996) and dynamic zone behavior.
- Early version of secure updates of dynamic zones Internet standard. This was previously available through named4. Named4 and named8 become obsolete and are replaced by “named” (new daemon), which provides named4 and named8 function.
- Since AIX Version 4.3.1, Dynamic Host Configuration Protocol (DHCP) is now a threaded implementation allowing improved scalability.

Supported Standards

The following standard RFCs and drafts are supported for IPv6:

- RFC 1883 — IPv6 Specification
- RFC 1884 — IPv6 Addressing Architecture
- RFC 1885 — Internet Control Message Protocol (ICMPv6) for the IPv6
- RFC 1886 — DNS Extensions to Support IPv6
- RFC 1887 — An Architecture for IPv6 Unicast Address Allocation
- RFC 1970 — Neighbor Discovery for IPv6
- RFC 1971 — IPv6 Stateless Address Autoconfiguration
- RFC 1972 — A Method for the Transmission of IPv6 Packets Over Ethernet Networks
- RFC 1981 — Path MTU Discovery for IPv6
- RFC 2133 — Basic Socket Interface Extensions for IPv6 Supported Draft:

TCP/IP Resolver Enhancements

New Dynamic Load Name Resolver APIs included in AIX Version 4.3.3 allow users to define their own Name Resolver module in addition to the existing Name Resolver methods, such as local, DNS/Bind, NIS, and NIS+. Users can

create a module, which may have one or all of the following five map types: Services, protocols, hosts, networks, and netgroup.

TCP Checksum Offload for Gigabit Ethernet

In the new gigabit Ethernet device driver, the workload of TCP checksum processing is offloaded from the AIX TCP/IP protocol stack to the adapter. This reduces the amount of CPU time spent computing checksums in the main CPU, therefore, allowing more packets to be processed. Along with the network buffer cache, this feature allows sending data with no data touching, therefore, helping to increase performance. This function benefits applications that transmit large blocks.

TCP/IP Socks Library

The AIX SOCKS API allows generic TCP/IP applications to connect to hosts through a generic TCP/IP proxy using Version 5 of the SOCKS protocol. Any application that only makes outgoing TCP connections can take advantage of this API without any code modification because the library will automatically handle the tunnel creation with a configured SOCKS5 server. Furthermore, the network administrators can configure the API to accept and route certain types of clients' requests to different servers, thus balancing the workload across multiple SOCKS5 servers. This library enables network administrators to allow limited access to external sites while maintaining network boundary controls.

IP Security on IPv4 and IPv6

New IP security protocols, Authentication Header (AH) for integrity and authentication and Encapsulating Security Payload (ESP) for confidentiality, are supported. IPsec implementation interoperates with other IPsec-enabled products, such as IBM Firewall Version 3.1 for AIX.

A key management mechanism is also supported since keys must be generated and distributed for IP security. A shared master key is manually distributed between systems at the tunnel endpoints, where one system generates and exports a master key file, and the other system uses this key file to import the master key definition. Session key protocols, both static for tunnel duration and dynamic, are also provided.

IP Security for v4 and v6 in AIX Version 4.3.1 was enhanced as follows:

- Addition of triple DES encapsulation (included in AIX Version 4.3.1 Bonus Pack).
- Increased options in pairing of authentication and encapsulation algorithms.

- Performance improvements in both Hashed Message Authentication Code - Secure Hash Algorithm 1 (HMAC-SHA1) and DES for PowerPC-based platforms.
- Crypto extensions can be dynamically loaded and unloaded.
- Improved logging of tunnel and filter events.

Several combinations of transforms and algorithms can be used to maximize interoperability and provide various levels of security:

- The RFC 1826 AH and the RFC 1829 ESP can be used either together or alone. The Keyed MD5 algorithm is available for RFC 1826 AH, and the DES CBC 4, DES CBC 8, and CDMF algorithms can be used with RFC 1829 ESP.
- The new 96-bit HMAC AH format and the new combined ESP with Authentication format can be used. These formats currently exist as IETF drafts. The HMAC MD5 and HMAC SHA1 algorithms can be used with this new AH format, and the DES CBC 4, DES CBC 8, CDMF, and the DES CBC MD5 combination algorithm can be used with the new ESP format. Replay protection can also be used with either of these new formats.

Note

Due to various countries' restrictions, DES and CDMF are separately installed from separate media. Customers who require the IPsec Confidentiality must install the separate components from separate media AIX Version 4.3 Bonus Pack.

Supported standards:

- RFC 1825 — Security Architecture for the Internet Protocol
- RFC 1826 — IP Authentication Header
- RFC 1827 — IP Encapsulating Security Payload (ESP)
- RFC 1828 — IP Authentication Using Keyed MD5
- RFC 1829 — The ESP DES-CBC Transform
- RFC 2104 — HMAC: Keyed-Hashing for Message Authentication
- RFC 2085 — HMAC-MD5 IP Authentication with Replay Prevention

TCP/IP Performance Enhancements

TCP/IP network interfaces support the thread mode on SMP systems, which:

- Allows IP incoming packets to be queued and picked up by threads, thus shortening interrupt time.

- When turned on, allows the processing of incoming IP packets to be shared among multiple processors.
- Works best with high bandwidth adapters, such as 1 Gbps Ethernet or ATM 155 Mbps.
- In streaming scenarios, can help increase the throughput of high-speed network interfaces on the receiving machine.

TCP has been changed to inherit TCP/IP socket options on new connections between TCP/IP clients and servers. This enhancement allows server programs to set the TCP/IP level socket options once on the listening socket and avoid setting options on the sockets associated with new connections, thus reducing path length (instruction) overhead.

Path Maximum Transmission Unit (PMTU) discovery is now the default. This feature helps AIX systems determine the best way to deliver maximum packet without fragmentation. As a result, it helps optimize throughput and increase network performance across any Internet affected by the packet size used.

TCP Selective Acknowledgements (SACK) helps TCP recover from multiple packet losses within the TCP window. It includes the following features:

- When SACK is enabled, TCP attaches SACK options to TCP acknowledgements, informing the sender that data has been received, but cannot be acknowledged by TCP due to missing data.
- The sending TCP can then retransmit only the missing data segments, thus reducing the number of packets transmitted and allowing TCP to recover more quickly than traditional retransmission methods.
- With SACK, TCP can also avoid dropping back into slow start when packets are lost, allowing TCP to continue to send data at near network capacity.

TCP/IP Gratuitous Address Resolution Protocol (ARP)

Gratuitous ARP provides these features:

- When the network adapter hardware address for an IP host changes, then all other hosts on the same physical network that have an ARP entry with the old hardware address update their entry with the new hardware address. The address update will take place during a non-intensive performance period, such as network interface bringup; therefore, it helps reduce the waiting time when data is ready to transfer through a TCP/IP connection.
- When assigning an IP address to an interface, AIX will detect whether another host is already configured with the same IP address by sending

an ARP request for the desired address. If the desired address is discovered to be already in use, the system will log this error to notify the system administrator.

14.1.7.7 Quality of Service

Advanced Internet services are based on Quality of Service (QoS) and policy-based networking. QoS embodies the idea of preferential treatment for certain traffic flows and enables service providers to implement and effectively manage service level agreements.

The AIX QoS implementation is based on the Internet Engineering Task Force (IETF) standards, Integrated Services (IntServ), and Differentiated Services (DiffServ). IntServ utilizes the Resource ReSerVation Protocol (RSVP) available to applications using the RSVP API (RAPI). DiffServ support includes IP packet marking for IP packets selected through filtering. The AIX QoS also offers bandwidth management functions, such as Traffic Shaping and Policing. Policy-based networking is an industry effort to define and standardize various network configuration parameters needed for effective network management. The AIX QoS scope covers both QoS and policy-based networking.

This enhancement to AIX provides system administrators with the benefits of both QoS support and policy-based networking in meeting the challenges of QoS offerings across complex networks.

14.1.7.8 Network File Cache

A `send_file` API with in-kernel Network File Cache is added for high-performance, file-to-network I/O. The buffers for this cache are allocated from the network memory allocator, which is expanded to 1 GB in AIX Version 4.3.2. Using this capability, file contents are cached in the Network File Cache for a subsequent re-use in a `send_file` operation. If the data is not present in the network file cache, the file is read from the file system and transmitted over the network. This function helps network servers (Web servers, file servers) improve I/O efficiency, thus improving the capacity in terms of operations/second.

Industry multithreaded network server application developers will appreciate that AIX Version 4.3.2 overcomes the *thundering herd* problem of wasted CPU cycles resulting from waking up multiple threads of a network server when new connections arrive. In AIX Version 4.3.2, only a single thread is awakened, therefore, reducing CPU overhead.

14.1.7.9 IP Addresses

There is no restriction on the number of IP addresses you can have per adapter. AIX does not restrict the address quantity. However, the system administrator may want to consider performance issues when assigning ip addresses to the adapter.

14.1.7.10 HTTP GET Engine

The new kernel extension HTTP Get Engine helps to increase the performance of Web serving by serving Web pages from AIX Network File Cache. Web pages are cached in Network File Cache. Using Power and PowerPC private segments, this cache can be extremely large and can be constrained by the physical memory in a system.

The AIX HTTP Get Kernel extension intercepts incoming HTTP get requests from the network and serves Web pages if found in the Network Cache without the overhead of get requests sent to user space Web server. If Web pages are not found in AIX Network File Cache, the HTTP Get request is queued to user space Web server. Furthermore, the Network File Cache Web page data is transmitted by reference without requiring any data copies.

This function, in conjunction with TCP/IP Checksum Offload capability in ATM 155 Mbps adapter and 1 GB Ethernet adapter, makes up the non data touching (by CPU) architecture. The kernel HTTP Get Engine, along with Network File Cache, helps reduce the number of instructions needed to process a get request dramatically, therefore, increasing server throughput. This capability is used by HTTP Server (powered by Apache) Web server software Version 1.3.6.

14.1.7.11 EtherChannel

Cisco System's network bandwidth aggregation and load balancing technology, called EtherChannel, builds upon standard and 802.3 Fast Ethernet and provides the function to aggregate a bandwidth of multiple Ethernet interfaces.

IBMs implementation for AIX of Cisco's EtherChannel defines a logical interface or channel, which can consist of two to four physical interfaces (or adapters). This logical interface or channel looks like an Ethernet interface to the upper layers. Any upper layer (IP, SNA, DLPI, for example) that can connect to an Ethernet adapter through AIX Common Data Link Interface (CDLI) network services is designed to work over an EtherChannel without any code change.

AIX's EtherChannel implementation is designed to work with Ethernet Switches that are compatible with Cisco's EtherChannel technology. Traffic

sent to the channel is sent on the network over one of the devices that is part of the channel. For IP traffic, the outgoing device is selected by hashing the destination IP addresses. For all other upper layer traffic, the outgoing device is selected by hashing the destination MAC addresses. Hashing of addresses helps ensure that traffic between a particular source and destination uses the same physical interface or adapter. This is used to avoid packets reaching the destination out of order, as some higher layer protocols have problems handling out-of-order packets. Packets received from any of the interfaces are sent to the higher layers for processing, regardless of which device they are received on.

14.1.7.12 DHCP Functional Upgrade

The Dynamic Host Configuration Protocol (DHCP) product has been enhanced at AIX Version 4.3.3 with a programming interface to the DHCP server for user extension or third-party software integration. Performance enhancements were made to the `nsupdate4` utility.

At the AIX 4.3.1 level, the DHCP updates were:

- Performance enhancement through data organization and threading
- Multiple database backend support through dynamic loading
- Out of protocol modifications to databases using an administration interface
- Logging level meaning and accuracy improvements
- Classless interdomain routing (CIDR) addressing
- Support for virtual subnets

14.1.7.13 NFS Version 3 Support

NFS Version 3 is supported using the `fs_max_read_size` command. It sets the maximum and preferred read size for NFS Version 3 servers. The default size is 32,678 bytes. The maximum is 65,536 bytes and the minimum is 512 bytes. NFS clients mounting after `fs_max_read_size` is set must use its set value. You cannot use this option to change the size existing mounts.

Using `fs_max_read_size` may require you to reduce the Version 3 read/write sizes when the mounts cannot be manipulated directly in the clients, especially during NIM installations on networks dropping packets with default read/write sizes of 32 KB. In this case, set the maximum size to a small value.

14.1.7.14 WebFS

AIX provides NFS server capability for WebNFS. Defined by Sun Microsystems, WebNFS is a simple extension of the NFS protocol that allows easier access to servers and clients through Internet fire walls.

A WebNFS-enhanced Web browser can use an NFS universal resource locator (URL) to access data directly from the server. An example NFS URL is:

```
nfs://www.YourCompany.com/
```

WebNFS works in tandem with existing Web-based protocols to provide data to clients. WebNFS also takes advantage of the scalability of NFS servers.

14.1.7.15 PC Services - CIFS Services

From PC Services, you can perform tasks associated with defined PC services or IBM AIX Fast Connect servers, which provide file and print services to networked personal computers that use System Message Block (SMB) or Common Internet File System (CIFS) protocols.

14.1.7.16 MPOA on ATM Enhancement

AIX 4.3.3 includes enhancements for multiple-protocol over ATM (MPOA) which provides improved management and performance of an ATM LAN Emulation network by combining multiple edge routers into a single router image. Device specific configurations are minimized with auto-discovery and device discovery protocol, while data paths are reduced from many hops between routers to a single hop between end clients. MPOA supports both standard Ethernet and 802.3 Ethernet in AIX 4.3.3.

14.1.8 System Performance Enhancements

AIX Version 4.3.1 introduced some very detailed performance tuning. Individual installations may see varying performance improvements based on the actual configuration of their systems. AIX Version 4.3.1 improved performance in the following ways:

- Streamlined the JFS locking to allow read access up to four times faster in an application where two processes sequentially read different parts of the same file.
- Improved Java performance up to 40 percent with new versions of the Java toolkit (1.1.4) and the JIT compiler and by turning on the compiler by default.
- IP Security Data Encryption Standard (DES) and Commercial Data Masking Facility (CDMF) encryption and decryption algorithms were

improved to take advantage of the PowerPC architecture. This results in an up to 300 percent performance improvement (up to four times the original throughput) on a 100 Mbps Ethernet.

- For the slower 10 Mbps Ethernet, the throughput is now at media speed but with a 50 percent lower CPU utilization. Similar performance gains are achieved on other media types, such as token ring.
- Optional fork interface primarily for use by applications, such as Web servers that require a lightweight fork behavior. The interface should be used for multithreaded applications that create new programs using the `exec()` call immediately after they call `fork()`.

14.1.9 System Security

The first section will discuss AIX Version 4.3.3 new directory exploitation of AIX users and groups utility. Following this, the Secureway Directory, AIX's LDAP, will be discussed.

14.1.9.1 Directory Exploitation

AIX Version 4.3.3 offers new directory exploitation of AIX users and groups. It provides a facility, which allows AIX users and group information to be optionally stored, replicated, and retrieved in an SecureWay Directory for fast access (local or remote), expandability, and reliability. When an AIX system is configured, user and group-related queries are sent to and responses are received from the SecureWay Directory. All AIX user data is securely stored, replicated, and managed by the SecureWay Directory server. For a collection of AIX systems that need to share a common view of user security information, this function can significantly reduce the number of administrative operations.

14.1.9.2 SecureWay Directory Version 3.1.1

SecureWay Directory is an open, cross-platform server optimized to support Lightweight Directory Access Protocol (LDAP)-enabled applications that integrate enterprise systems. Providing a unified architecture that allows users to share data with people, applications, and network resources, SecureWay Directory helps improve communication, speed development, and deployment of Web applications and increases the security of the network. Utilizing the power of DB2, Universal Database (UDB) and its transactional data store, the directory extends the performance and availability of DB2 to an enterprise directory service.

Note

You may only use the DB2 UDB component in association with your licensed use of the SecureWay Directory.

The SecureWay Directory, previously announced as eNetwork LDAP Directory, has been rebranded and renamed under the SecureWay brand to more closely align with the IBM eBusiness portfolio for SecureWay Software. These products provide an integrated solution and a secure network platform for our customers to implement an e-business.

The new version of the SecureWay Directory steps up to the Internet Engineering Task Force (IETF) LDAP V3 support based on RFC 2251, 2252, 2253, 2254, and 2256. Many new features are provided over the eNetwork LDAP Directory, which was based on LDAP V2. LDAP V3, provides enhancements to both the LDAP protocol and the supported schema.

The new LDAP V3 protocol features include:

- Referrals — A list of server URL addresses are returned to a client whose request cannot be serviced. The client can use the server locations returned to continue the operation.
- Controls — Extension information can be added to a request for an LDAP operation.
- Extended operation plugin support — Additional operations can be defined for services not available elsewhere in the V3 protocol. Clients can request and receive responses with predefined syntaxes and semantics.

The SecureWay Directory server uses attribute-type definitions, object-class definitions, and other information called schema to determine how to match a filter or attribute value against an attribute of a directory entry. The schema matching also determines whether or not add or modify operations are permitted. The breadth of supported schema definitions has grown to support, not only the schema defined by LDAP V3, but also encompasses IBM common schema and Directory-Enabled Networks (DEN) schema.

Subclassing enables new object classes to be defined that inherit the object class definitions and attributes of its parent class. The new object may be defined with additional or changed attributes. Schema update operations are checked against the schema class hierarchy for consistency before being processed. Additionally, the Directory permits authorized users to dynamically define new attributes and object classes to enhance the pre-defined directory schema.

SecureWay Directory has provided a migration utility to convert your eNetwork LDAP Directory V2.1 schema definitions to LDAP V3 format. No migration is required of the directory data. eNetwork LDAP Directory V2.1 data will work with the SecureWay Directory V3.1.1 server.

Server-specific information, Directory System Agent (DSA)-Specific Entry (DSE), is contained in a read-only repository, RootDSE, that contains the following information:

- Suffixes supported by the local directory server
- Distinguished Name (DN) of the subschema entries known by the server
- List of alternative (replica) servers
- LDAP version implemented by the server
- List of supported extended operations
- List of supported controls
- List of supported Simple Authentication and Security Layer (SASL) security features
- Server configuration information

SASL, defined in RFC 2222, is a framework for adding pluggable authentication support for connection-based protocols. The directory server invokes the SASL plugin functions to perform authentication following a bind request from a client and returns the results to the client.

Two methods of authentication are supported:

- Challenge/Response Authentication Mechanism —Message Digest 5 (CRAM-MD5)
- Secure Socket Layer (SSL)

Several new features have been added to address security:

- In addition to certificate authentication for the server, which was available with Version 2.1, the Directory now supports SSL client certificate authentication based on public keys, which provides the means for setting up a protected communication channel between the client and server. A user with a public key certificate signed by a Certificate Authority can use the certificate to authenticate to the directory server.
- Full SSL Java Naming and Directory Interface (JNDI) support.
- Encryption of passwords in the Directory prevents passwords from being compromised via database queries or file lookup.

A Change Log has been included in this release, which logs add, delete, and modify operations to the directory server and changes to the change log itself. A client can access the Change Log and update its own replicated copy of the directory data by applying the changes.

Data can be stored, retrieved, and managed in the Directory using a native language code page for either single-byte or double-byte languages. Data is converted to the Universal Code Set (UCS) Transformation Format (UTF-8) character strings before being sent to and from the server. The data can be stored as either UTF-8 or as local codepage strings depending on the database configuration. This version has translated messages for Group 1 languages and Czech, Polish, Hungarian, Russian, Catalan, and Slovakian.

SecureWay Directory can be administered and configured from a Web browser-based GUI. The administrator can:

- Perform initial setup of the directory
- Change configuration options
- Manage the daily operations of the directory

User access control is provided for information stored in the Directory and can be defined by an administrator. From a Web browser, users can search for, or add to, information in the Directory. In addition, the Java-based Directory Management Tool is provided to allow a user to perform these tasks:

- Connect to one or many directory servers via secure or unsecure network connects
- Browse the directory tree or directory schema
- Add, edit, modify, and delete objects, object classes, and attributes in the directory

Client access to the SecureWay Directory is supported using LDAP or HTTP protocols. AIX client applications can be developed using the enhanced elements provided for supporting LDAP V3 protocols and APIs. These elements are provided by the SecureWay Client SDK, which consists of:

- Client libraries that provide a set of C-language APIs
- C header files
- Documentation (in the form of HTML files)
- Sample programs
- Executable versions of the sample programs

Additionally, the following components are provided for developing Java applications that use Sun's JNDI. This permits Java applications to access LDAP-compliant directory servers:

- JNDI class files
- A set of class files for the LDAP service provider
- Documentation

The LDAP libraries and utilities provided with the SDK utilize the SSL libraries, if present. The SSL libraries are provided as part of the IBM Global Security Kit (GSKit). When GSKit is installed, the LDAP library will dynamically load the SSL libraries and use them to enable support of SSL. The LDAP library is fully functional regardless of the presence of SSL. GSKit Version 3.0.1 is available on the AIX Bonus Pack 4.3.

The U.S. government's regulations regarding the export of SDKs, which provide support for strong encryption, continue to evolve. This has resulted in changes in the way IBM packages the SecureWay Directory Client SDK and the manner in which LDAP applications gain access to the strongest SSL encryption algorithms (which include 128-bit and triple DES encryption). The point of control, with respect to available levels of encryption, is now the application.

Any LDAP application that uses the SecureWay Directory Client SDK Version 3.1.1 with the required level of GSKit 3.0.1.84 (or higher) has default access to 56-bit DES encryption (over SSL). This is the case for LDAP applications (new and existing) that use either the domestic or general export versions of SecureWay Client SDK Version 3.1.1.

For an LDAP application to access the stronger SSL cryptographic encryption algorithms, the application must use a new function that sets the cipher support to 128-bit/triple-DES and registers the application for the stronger cryptographic encryption algorithms. Without this function, LDAP applications have default access to a maximum of 56-bit DES encryption for SSL connections. To invoke the new function, your application must be linked with the appropriate static library that exports it. These static libraries are distributed via the IBM SecureWay Directory Security Enabler V3.1.1 package (5648-D14). For users within the U.S. and Canada, this package can be download from URL:

<http://www.software.ibm.com/network/directory>

These static libraries, which provide unrestricted cipher support and applications developed with these libraries, may be exported outside the U.S.

and Canada only with the appropriate export license, as provided by the U.S. government.

The SecureWay Directory is Tivoli-ready.

14.1.10 Network Security

The following sections discuss some of the network security functions supported by AIX Version 4.3.

14.1.10.1 IP Security Enhancements

The following outlines the TCP/IP Security Enhancement at the various AIX 4.3 levels:

- Remote TCP/IP Commands

AIX Version 4.3.1 offers the following secure remote TCP/IP commands: `rsh`, `rcp`, `rlogin`, `telnet`, and `ftp`. With this capability, Kerberos 5 authentication is used between these commands and server daemons, therefore, avoiding the need for user passwords to pass in the clear on the network. Instead, Kerberos 5 credentials are used to authenticate users. User credentials can be forwarded to the server.

- Tracing IP Security and IKE Messages

Logic has been added to allow better tracing of the IP Security and IKE messages. The output of logging is now formatted in readable format and AIX auditing has been implemented. Users can now pinpoint configuration failures and view audit logs to determine security attacks.

- On-Demand Tunneling

Dynamic tunnels only need to be defined one time and then will be activated only when traffic matching the criteria set out in the policy was sent or received. This feature is beneficial to users because the functions for negotiating, computing, and refreshing session keys will only be performed when necessary.

- Filtering Based on IP Address Ranges

IKE Tunnels can be created that specify a range of IP addresses (starting and ending IP address range endpoints) that allow tunneling over multiple IP addresses. It allows users to easily define tunnels for ranges of addresses.

- Web-based System Manager Interface for Filters

A GUI-based tool can now be used to configure and manage manual tunnels, static and dynamic filter rules, and importing and exporting tunnel

definitions. The tool is now consistent across IP Security and is NLS-enabled; however, SMIT is not available for this option.

- **Certificate-based Use of Digital Signatures for IKE Authentication**
IKE tunnels (dynamically negotiated secure tunnels) have been enhanced to use digital certificates for authentication. Authentication is accomplished by signing IKE messages using X.509 certificates. Certificates may be stored locally. This enhancement enables the deployment of Virtual Private networks with a large number of endpoints. Such network configurations may present a savings in cost and administration of typical leased line installations. Support for certificate-based use of digital signatures for IKE Authentications is planned to be available November 5, 1999 for new orders and with APAR IY02769 for existing customers.

14.1.11 Performance Analysis Tools Enhancements

To provide more system performance information and to improve system serviceability, AIX Version 4.3.0 contains the following tools, which formerly were in the Performance Toolbox licensed program product. These tools normally require root authority and should only be used by experienced IS professionals for system performance analysis.

These character-based tools provide deeper statistical views into local AIX system performance metrics. The majority of these tools have been updated to work with both 32-bit and 64-bit systems and applications (exceptions are noted).

- **bf** — Traces the memory access patterns of running programs (32-bit applications only)
- **fdpr** — Restructures programs based on observed execution
- **filemon** — Traces detailed disk I/O activity
- **fileplace** — Displays file block placement within a logical or physical volume
- **lockstat** — Reports of contention for locks
- **netpmon** — Traces detailed network I/O activity
- **rmss** — Simulates smaller real memory to assess the memory requirements of programs
- **stem** — Traces the control flow and CPU use of programs without access to their source code. This was for 32-bit applications only until AIX 4.3.1 when stem was enabled for 64-bit systems

- `svmon` — Reports the current use of real and virtual memory
- `tprof` — Reports the CPU use of individual programs, subroutines, and system components

Performance Analysis Tools in AIX Version 4.3.3 add the following improvements:

- `fdpr` — The usability and reliability of this tool has been improved by adding automatic selection of options based on the type of execution used.
- `gennames` — This tool is used in conjunction with the new off-line mode capability found in `tprof`. It consolidates and simplifies the information needed by `tprof` to process symbol, loader, and extension information.
- `svmon` — This tool has been enhanced with usability, scalability, and speed improvements on RS/6000 Enterprise Server Systems. In addition, it supports the new Workload Management function offered in AIX Version 4.3.3.
- `tprof` — This tool has been enhanced to produce separate statistics per processor when running off-line and in conjunction with a new trace option to produce one trace log file per processor.

In addition, the Performance Analysis Tools adds the following new tools:

- `ipfilter` — This tool sorts the information provided by the `ipreport` command and presents it in table format. The `ipfilter` allows the user to select which operation headers (NFS, UDP, IPX, and ICMP) to view.
- `pprof` — A lightweight, trace-based tool, which collects a system's process and thread information. Reports are generated in several formats, including a family view. The family view displays all parent-child relationships for all processes and threads. This tool is especially helpful in pinpointing system degradation when caused by multiple processes.
- `topas` — This tool utilizes the Performance Toolbox System Performance Measurement Interface to sample and report a wide variety of local system statistics, which include event, file, disk, memory, network, paging, process, and queue information. This tool provides an easy-to-use snapshot of overall system activity.

The Performance Toolbox Agent (PTX) has been upgraded with the ability to process performance metrics by activity rather than by a fixed name. This capability allows system metrics to be logged only when customers' specified thresholds are exceeded. Another improvement to PTX is the ability to create a file, which contains shell commands to be executed when recording files

are deleted. This feature allows users to merge, rename, or move recording files automatically. Finally, the PTX agent has been modified to process a larger number of local processes on each monitored system.

14.1.12 PC Interoperability - AIX Fast Connect

This section will describe the features of AIX Fast Connect.

AIX Fast Connect Release 2.1.1 for Windows and OS/2 includes integration with DCE/DFS. This feature allows user authentication with the DCE security server. DFS directories can be shared with (exported to) PC clients. PC client access is controlled by the login context acquired as a result of DCE authentication. Fast Connect offers DCE authentication and DFS access without requiring DCE/DFS client software to be installed on each of the PC clients requiring DCE/DFS access. This allows centralized user/resource management for PC clients using DCE/DFS, avoiding cost, complexity, and extensive management overhead associated with installing DCE/DFS client on every desktop.

AIX Fast Connect supports Windows 95/98/2000 logon. Windows NT clients are supported using Network Client for Windows NT. This feature of AIX Fast Connect allows support of home directories, startup scripts, and roaming user profiles for PC clients without requiring Windows NT domain controllers.

With the AIX Fast Connect for Windows feature, AIX file and print services are available to Windows 95/98 and Windows 2000 users as part of their Network Neighborhood -- no additional software is required. Windows clients can access AIX files using the Journaled File System (JFS), CD File System (CDFFS), Network File System (NFS) mounted subsystems, and AIX printing services using Microsoft's Common Internet File System (CIFS) and Server Message Block (SMB) protocols over TCP/IP. It supports CIFS and SMB protocols on NetBIOS-over-IP transport and interoperates with Windows NT Server.

AIX Fast Connect logon server is no longer required to be located in the same IP subnet as Windows NT clients.

Windows NT client support requires use of encrypted passwords and AIX Fast Connect logon server located in the same IP subnet as clients. IBM Network client software for Windows NT and Windows 95 is available for download at the following IBM Web sites:

http://service.boulder.ibm.com/asd-bin/doc/en_us/winntc12/f-feat.htm

http://service.boulder.ibm.com/asd-bin/doc/en_us/win95c1/f-feat.htm

AIX Fast Connect Release 2.1.1 is shipped as a billable AIX feature.

It has a rich set of functions, such as:

- AIX File System Support with Common File Names
- AIX Print Queue Support
- Local Master Browser Support
- WINS Server and Client Support
- Opportunistic Locks
- NT Domain Passthru
- AIX Unicode support
- Command line, SMIT, Web-based systems manager
- Encrypted Passwords
- DCE/DFS File Systems Support
- Netlogon and Profile Support AIX ACLs support

Some of the key attributes are:

- AIX Integration
 - AIX installation, tracing, logging, auditing, and fix services
 - Exploits AIX's core services for efficiency, scalability, and reliability
- Performance
 - Better throughput than other commercial AIX-based File/Print servers
- Management and Administration
 - SMIT and Web-based System Manager interfaces
- Affordable
 - AIX feature with modest one time charge
- Service and Support
 - IBM Worldwide

Additional product information can be found at:

www.ibm.com/servers/aix/products/ibmsw/manage/fastconn.html

Note

Fast Connect does not support Netware Clients. In this case, a non-IBM product, such as Total Advanced Server from Syntax, could be recommended.

14.1.13 Java Support

AIX Version 4.3 includes IBM's port to AIX of Sun's Java Developer's Kit (JDK) Version 1.1.6. The IBM Just-in-Time (JIT) Compiler 3.0, included with JDK 1.1.6, compiles Java byte code to native machine code at runtime to boost Java performance up to 25 times. Since Java technology is rapidly changing, installation of JDK has been modified to accommodate multiple versions of JDK in AIX Version 4.3 to enable Java-based applications, which require different versions of JDK. This allows you the flexibility to take advantage of Java technology while maintaining a stable JDK version.

Java Development Toolkit (JDK) Version 1.1.6 with the IBM Just-In-Time (JIT) Compiler Version 3.0 also include Euro currency symbol support and performance enhancements. Enhancements are made in the base operating system to improve Java performance, scalability, and robustness.

The following performance improvements for Abstract Windowing Toolkit (AWT) are included in this release:

- Software to improve the performance of drawing and filling rounded rectangles
- Software to improve colormap initialization

Applications and applets do not have to change to obtain these AWT performance improvements.

The AIX 4.3 Java Port for AIX includes:

- Appletviewer
- Optimized Java interpreter
- Java class compiler
- Source-level debugger
- Java runtime interpreter
- Java classes, which include the following APIs:
 - Java Database Connectivity (JDBC), allowing database access from Java

- JavaBeans as a component model level interface
- Remote Method Invocation (RMI)
- Java Security API, allowing both low-level and high-level security function to be built into Java applications
- Tool for automatic generation of HTML documentation for an applet
- Runtime libraries for Java multimedia links

14.1.13.1 Java Security

AIX Developer Kit, Java Technology Edition, Version 1.1.8, delivered on AIX Version 4.3.3, provides the following security enhancement:

- Java Security Migration Aid — The Security Migration Aid provides the more robust, policy-based security model of the Java 2 platform in the Java 1 platform environment. The Security Migration Aid is intended to help users migrate from the relatively simple System Security Java 1 platform security mode to the finer grained Java 2 platform model. The Security Migration Aid supports the use of security policies, permissions, tools, and Java runtime security managers as defined in the Java 2 platform security model.

AIX Developer Kit, Java Technology Edition, Version 1.1.8, delivered on AIX Version 4.3.3, provides the following features:

- Java Remote Method Invocation — Internet Inter-ORB Protocol (RMI-IIOP)

A new version of RMI that runs over IIOP and interoperates with Common Object Request Broker Architecture (CORBA) Object Request Broker (ORB) and CORBA objects programmed in other languages. This new feature combines the simplicity of Java RMI programming and the heterogeneous interoperability of CORBA. The RMI-IIOP toolkit includes code generators that work directly with Java RMI class files used by applications on the client side and code that wraps object implementations on the server side. In this way, the developer can work in 100 percent Pure Java on both sides and does not have to know the CORBA IDL (Interface Definition Language). This enables Java client applications to access other server objects besides those written in Java.

- Java Database Connectivities (JDBC) — Object Database Connectivities (ODBC) Bridge

The Bridge enhances the capability of enterprise customers to communicate with databases via Java. The Bridge provides JDBC access to databases with ODBC drivers.

- Swing 1.1.0

This version is part of the Java Foundation Classes (JFC) that implements a new set of GUI components with a pluggable look and feel. Swing is implemented in Java and is based on the Java 1 platform Lightweight UI Framework. The pluggable look and feel lets you design a single set of GUI components that can automatically have the look and feel of any OS platform (Windows, Solaris, Macintosh). Swing components include both 100 percent Pure Java versions of the existing AWT component set (button, scrollbar, and label).

- Big Decimal

IBM has enhanced Java's Big Decimal math class by adding support for floating point arithmetic. The IBM Big Decimal class implements the decimal arithmetic defined in the ANSI standard X3.274-1996. The advantages are:

- A full-function decimal floating point arithmetic. For example, mantissa length information is not lost; so, trailing zeros can be preserved (for example, $1.2 \times 2 = 2.40$, not 2.4)
- Exact results (for example, $0.9/10 = .09$, not .089999996).
- The precision of the arithmetic is freely selectable by the programmer, not limited to a choice from one or two alternatives.
- Robust arithmetic operations — There is no wrap of integers at certain sizes. Ill-defined or out-of-range results throw exceptions.
- A single class can be used for all decimal numbers. There is no arbitrary distinction between integers and floating point numbers.
- Both scientific (where one digit is shown before the decimal point) and engineering (where the exponent is a multiple of three) exponential notations are supported.
- Exponents in the range E-999999999 through E+999999999 are supported.

14.1.14 Additional Printer Support

AIX Version 4.3.0 supports the following printers:

- Bull Compuprint PageMaster 1015
- Bull Compuprint PageMaster 1021
- Bull Compuprint PageMaster 1025
- Bull Compuprint 1070
- Bull Compuprint PageMaster 1625
- Bull Compuprint PageMaster 200

- Bull Compuprint PageMaster 201
- Bull Compuprint PageMaster 411
- Bull Compuprint PageMaster 413
- Bull Compuprint PageMaster 422
- Bull Compuprint 4/51
- Bull Compuprint 4/54
- Bull Compuprint PageMaster 721
- Bull Compuprint PageMaster 815
- Bull Compuprint PageMaster 825
- Bull Compuprint 914
- Bull Compuprint 914 N
- Bull Compuprint 922
- Bull Compuprint 923
- Bull Compuprint 924
- Bull Compuprint 924 N
- Bull Compuprint 956
- Bull Compuprint 970
- Bull PR-88
- Bull PR-88 VFU Handling
- Bull PR-90
- Canon Laser Shot LBP-B406/S/D/E,A404/E,A304E
- Canon LASER SHOT LBP-A404PS/Lite
- Canon LASER SHOT LBP-B406G
- Dataproducts BP2000 Line Printer
- Dataproducts LZR 2665 Laser Printer
- Hewlett-Packard LaserJet II
- Hewlett-Packard LaserJet III
- Hewlett-Packard LaserJet IIISi
- Hewlett-Packard LaserJet 4,4M
- Hewlett-Packard LaserJet 4 Plus/4M Plus
- Hewlett-Packard LaserJet 4Si
- Hewlett-Packard LaserJet 4V
- Hewlett-Packard LaserJet 5Si/5Si MX
- Hewlett-Packard LaserJet 5Si Mopier
- Hewlett-Packard Jet Direct Adapter
- Hewlett-Packard Direct EX
- Hewlett-Packard LaserJet Color
- IBM 2380 Personal Printer II
- IBM 2380 Plus printer (Model 2)
- IBM 2381 Personal Printer II
- IBM 2381 Plus printer (Model 2)
- IBM 2390 Personal Printer II
- IBM 2390 Plus printer (Model 2)

- IBM 2391 Personal Printer II
- IBM 2391 Plus printer (Model 2)
- IBM 3112 Page Printer
- IBM 3116 Page Printer
- IBM 3130 LaserPrinter
- IBM 3812 Model 2 Page Printer
- IBM 3816 Page Printer
- IBM 4019 LaserPrinter
- IBM 4028-NS1, AS1 Laser Printer
- IBM 4029 LaserPrinter
- IBM 4037 LP printer
- IBM 4039 LaserPrinter
- IBM 4070 IJ printer
- IBM 4072 ExecJet®
- IBM 4076 IJ printer
- IBM 4079 Color Jetprinter PS
- IBM 4201 Model 2 Proprinter® II
- IBM 4201 Model 3 Proprinter III
- IBM 4202 Model 2 Proprinter II XL
- IBM 4202 Model 3 Proprinter III XL
- IBM 4207 Model 2 Proprinter X24E
- IBM 4208 Model 2 Proprinter XL24E
- IBM 4208 Model 502 Proprinter XL24EK
- IBM 4212 Proprinter 24P
- IBM 4216 Personal Page Printer, Model 031
- IBM 4216 Model 510
- IBM 4224 Printer, Models 301,302,3C2,3E3
- IBM 4226 Printer
- IBM 4232-302 Dot Matrix Printer (when configured as an IBM 4202)
- IBM 4234 Dot Printer, Model 013
- IBM 4247 Line Matrix Printer
- IBM 4303 Printer
- IBM 4317 Printer
- IBM 5202 Quietwriter® III
- IBM 5204 Quickwriter®
- IBM 5327 Model 011
- IBM 5572 Model B02
- IBM 5573 Model H02
- IBM 5575 Model B02/F02
- IBM 5577 Model B02/F02/H02/G02/FU2/J02/K02
- IBM 5579 Model H02/K02
- IBM 5584 Model G02/H02
- IBM 5585 Model H01

- IBM 5587 Model G01
- IBM 5587 Model H01
- IBM 5589 Model H01
- IBM 6180 Color Plotter
- IBM 6182 Auto Feed Color Plotter
- IBM 6184 Color Plotter
- IBM 6185-1 Color Plotter
- IBM 6185-2 Color Plotter
- IBM 6186 Color Plotter
- IBM 6252 Impactwriter®
- IBM 6262 Printer
- IBM 6400 Printer
- IBM 7372 Color Plotter
- IBM Network Color Printer
- IBM Network Printer 12
- IBM Network Printer 17
- IBM Network Printer 24
- Lexmark 2380 Plus printer (Model 3)
- Lexmark 2381 Plus printer (Model 3)
- Lexmark 2390 Plus printer (Model 3)
- Lexmark 2391 Plus printer (Model 3)
- Lexmark Forms Printer 4227
- Lexmark 4039 plus LaserPrinter
- Lexmark ValueWriter 600
- Lexmark Optra laser printer
- Lexmark ExecJet IIc
- Lexmark 4079 Color Jetprinter Plus
- Lexmark Optra Plus laser printer
- Lexmark Optra C color laser printer
- Lexmark Optra E laser printer
- Lexmark Optra N laser printer
- Lexmark 4047 ValueWriter 600
- Lexmark 4076 ExecJet 11c Printer
- Lexmark Grand Prix
- OKI MICROLINE 801PS/+F,801PSII/+F,800PSIILT
- Printronix P9012 Line Printer
- QMS ColorScript 100, Model 20
- Texas Instruments OmniLaser 2115 Page Printer

AIX Version 4.3.1 adds support for the following printers:

- Hewlett-Packard LaserJet 4000
- InfoPrint 20

AIX Version 4.3.2 adds support for the following printers:

- Lexmark Optra Color 1200 printer
- Lexmark Optra Color 40 printer
- Lexmark Optra Color 45 printer
- Lexmark Optra K 1220 laser printer
- Hewlett-Packard LaserJet 8000
- Hewlett-Packard Color LaserJet 8500
- IBM InfoPrint 32

AIX Version 4.3.3 adds support for the following printers:

- Lexmark Optra W810 laser printer
- Lexmark Optra T laser printer family
- Lexmark Optra S plus laser printer family
- Lexmark Optra M410 laser printer
- Lexmark Optra E310 laser printer
- InfoPrint 40
- Hewlett-Packard 2500C Professional Series color printer
- Hewlett-Packard D640 printer
- Hewlett-Packard 8100 printer
- Hewlett-Packard Color LaserJet 4500 printer

AIX also has additional standard parallel printer support for CANON LIPS4 and EPSON ESC/P Japanese printer data stream.

European A3 paper sizes are supported for the HP printers previously listed.

AIX Version 4.3.2 can now view and manage six-digit print job numbers. Specifying the six-digit (wide) view helps to eliminate ambiguities with more than 1,000 print jobs in the system.

14.1.15 Support and Service

In this section, information about compatibility between AIX Versions 3 and 4 will be described. Planning information, such as packaging, ordering information, shipped documentation, automating software upgrades using Software Subscription, and the Terms and Conditions of licensing are discussed.

14.1.15.1 Compatibility between AIX Versions 3 and 4

All appropriately written AIX applications, written using AIX Version 3.2, or later, for POWER-, POWER2-, and PowerPC-based models, run on AIX Version 4 without recompilation for those same models. The only exceptions to this statement are applications using:

- Their own loadable kernel extensions
- Certain high-function terminal control interfaces
- X11R3 input device interfaces
- The LAN device driver interface
- SCSI device configuration methods (IHVs)
- The nlist() interface
- DCE threads
- Applications compiled using POWER2- or PowerPC-specific compiler options, but run on models other than POWER2 or PowerPC

Any program that must run in all environments — POWER, POWER2, and PowerPC (601 and above) — must be compiled using the common mode option of the compiler. Programs compiled to exploit POWER2 technology must be run on POWER2-based processors. Programs compiled to exploit PowerPC-based technology must be run on PowerPC-based processors. Existing code need not be recompiled to run.

Applications created on a system using AIX Version 4 may not function properly on a system using AIX Version 3.

Applications must have been created using the AIX shared libraries for these statements to apply.

AIX Version 3 and Version 4 Client and Server Compatibility

An RS/6000 system using AIX Version 3.2, or later, can operate as a server system for client machines using AIX Version 4, with the following exceptions:

- Network Install of Version 4 clients
- Service SNA or X.25 to Version 4 clients
- Service HCON to Version 4 clients
- Service CGE extensions of PEX and PEX-PHIGS
- Use Version 4 client install formats

An AIX system using AIX Version 4 can operate as a server system for client machines using AIX Version 3.2, or later, as long as the proper compatibility options are installed. All statements about binary compatibility apply in this case; Version 4 applications may not execute properly on Version 3 systems using remote network mounts of file systems.

In both cases, minor administrative changes are needed to the AIX Version 3 systems to support the new AIX Version 4 LFT terminal type.

14.1.15.2 Planning Information

A brief description of the packaging, unlimited user licensing of AIX 4.3.3, Software Subscription and the older Upgrade Protection Plan, and the licensing Terms and Conditions will follow.

Packaging

AIX Version 4.3.3 is distributed in one package with the following:

- IBM International Program License Agreement in a multi-language booklet and its License Information (LI)
- Proof of Entitlement (PoE)
- One CD-ROM

Unlimited User License

Since February 7, 2000, an unlimited user license is included when ordering most new RS/6000 systems. IBM eliminated additional user charges for the AIX operating system. RS/6000 systems will now come standard with an unlimited user license for the current release of AIX. Also, with this announcement, existing AIX V4 user licenses allow an unlimited number of users, and release upgrades from AIX 3.2, 4.1 and 4.2 to AIX 4.3 will now be for unlimited users. Customers upgrading AIX on their SP systems now use the same upgrade feature as other RS/6000 models, ordering one upgrade for each SP node on which the new level of AIX will be installed. Existing AIX Software Subscription customers are now entitled to unlimited users.

AIX Version 4.3 is available on all new RS/6000 systems shipping today.

While the base AIX operating environment is included at no charge with the base RS/6000 system, a separate AIX software order is required for registration purposes and to accommodate selectable options.

Ordering for AIX Fast Connect for Windows and AIX Fast Connect for OS/2, which was announced prior to this release, will stay the same.

14.1.15.3 Publications

The following additional hardcopy publications in Table 162 are shipped with AIX Version 4.3.3 or can be ordered from IBM after the planned availability date. To order, contact your IBM representative.

Table 162. AIX Publications

Title	Publication Number
<i>AIX Version 4.3 Installation Guide</i>	SC23-4112
<i>AIX Version 4.3 Network Installation Management Guide and Reference</i>	SC23-4113

Publications can now be viewed using the new Documentation Library GUI, which offers easier access to online documentation with a single integrated GUI that allows users to read, navigate, and search online HTML documentation.

14.1.15.4 Welcome Center for AIX 4.3

The theme of the AIX 4.3.0 Welcome Center was “RS/6000 on Mars” and contained a vibrant graphic theme, which revolves around a Mars space station.

The graphic theme of the AIX 4.3.2 RS/6000 Welcome Center is “GET OUT THERE...Find the Edge!”. Based on the American Motors General Hummer all terrain vehicle, the RS/6000 and AIX product family is portrayed as a rugged vehicle designed to *take you where you want to go*. Within the Welcome Center home pages, a fully textured mapped, 3D rendered, and animated model of the Hummer carries you on a journey through hostile but captivating and beautiful terrain in the comfort of this powerful four-wheel drive vehicle.

Functionally, the purpose of the Welcome center is to introduce new users to the new features and performance of their new RS/6000 system and to facilitate quick setup and configuration.

Similar to previous versions of the Welcome Center, this browser-based application contains a wealth of information about the user’s RS/6000 and AIX system. The application features link to more information on the World Wide Web, including:

- RS/6000 and AIX Product and Marketing Literature
- Local RS/6000 and AIX Web sites (geography-specific)
- RS/6000 Business Partners and Software Developers
- Application Development Tools

- RS/6000 Custom Solutions
- RS/6000 and AIX System Documentation
- IBM Service and Support Offerings
- RS/6000 FTP Sites and Newsgroups

In combination with the information provided locally on each RS/6000 system, the Welcome Center provides a wealth of customer information. The Welcome Center also contains a multimedia showcase with games, video content, and other audio-visual ways to demonstrate the performance and features of every RS/6000 and AIX system.

To demonstrate execution of IBMs Network Computing strategy within the RS/6000 product family, the Welcome Center is a Netscape Navigator-based application with a default *home page* and a large library of local HTML-format documents. Available immediately at system startup within the AIX boot sequence on pre-installed RS/6000 systems, the Welcome Center contains four primary chapters of information:

- About Your System

To familiarize the end user with the features of their RS/6000 system, this section provides an overview of the features of this system and Internet-based hyperlinks to a large library of more online IBM information.

- System Setup and Registration

To facilitate quick setup and system configuration, this section contains system setup information and hyperlinks to the new RS/6000 Customer Registration Web site.

- Just Imagine

This section contains a showcase of system features and performance organized around the top RS/6000 solution segments.

- Network Computing
- Decision Support
- Transaction Processing
- Visual Systems
- Technical Computing

Each of these sections contains information about each of these large and growing RS/6000 solution environments and specific information about RS/6000 products in each category.

- Contact IBM

To provide quick access to IBM RS/6000 and AIX worldwide service and support offerings, this section contains a directory of the many support services available to each of IBM's customers. This section also contains a directory of other valuable resources on the World Wide Web, such as online periodicals and application-porting assistance.

AIX Desktop Access and Uninstall Support

To accommodate easy access, the RS/6000 Welcome Center is available as a no-charge application on all pre-installed systems. The Welcome Center launches on these systems during the initial power-up sequence and is also available, after power-up, as an icon within the AIX Application Manager folder in the desktop environment.

The RS/6000 Welcome Center is easily uninstalled by selecting one of two optional uninstall features available within the exit page of this application. If the customer selects this uninstall option, the Welcome Center files are removed from the customer's system, making this space available for other applications. To preserve this function for other applications, the uninstall feature is designed not to uninstall the Netscape /Communicator Navigator or Ultimedia Services/6000.

Availability and Translation Support

The Welcome Center is currently distributed as a no-cost application within the standard AIX preload content from all RS/6000 manufacturing sites around the world. The RS/6000 Welcome Center in the order entry system is available as a selectable AIX option (default: YES) for all pre-installed RS/6000 systems. Since the Welcome Center content is heavily branded, non-IBM customers may elect to change the default setting to avoid receiving the Welcome Center on their systems. The Welcome Center is NLS-enabled. When the geographies provide translations, the Welcome Center pre-install image is updated to include other languages.

System Requirements

The RS/6000 Welcome Center is designed to install and operate on a minimally configured RS/6000 workstation or server with 32 MB of memory and a graphical display. The HTML source files, graphic images, and audio/video content for the Welcome Center consume approximately 80 MB of disk space. The Welcome Center requires Netscape Communicator 4.0.7 or later, Netscape Navigator 4.0 or later, and multimedia support, such as Ultimedia Services 2.2.1.0 or later, and speakers.

14.1.15.5 Software Subscription

A new enhanced program, Software Subscription for AIX (5692-SSO), is available for all new customers to more easily manage their program

currency and upgrade budget. Effective September 13, 1999, all new contracts for this offering will be under the new Software Subscription for AIX (5692-SSO).

Software Subscription for AIX provides RS/6000 customers with a streamlined process to keep their system's software current. The new, "no-signature" contract provides price protection and the convenience of a single, enterprise-wide administration point. All RS/6000 customers can now use the Web to sign-up for notification via e-mail of new AIX releases. After receiving notification, customers can also use the Web to order delivery of the new releases. This convenience, combined with automatic contract renewal, greatly simplifies software maintenance administration and helps improve expense management.

Software Subscription is extremely flexible. All of the software products you select, including products you add throughout the year, will be billed at an annual rate, therefore, allowing you to both maintain a single point of administration and predictably manage the cost of software. Periodic charge subscriptions are subject to a one-year minimum. The periodic charge will continue until the customer requests a discontinuance with a 30-day written notice.

This program is available with no minimum order requirements to all RS/6000 customers. Subscription offering for upgrades to AIX and related software products provides price protection and the convenience of a single, enterprise-wide administration point.

The price is determined by the products and product options selected.

Upgrade Protection Plan

The current AIX Upgrade Protection Plan (5692-ADV) will continue to provide upgrades for customers with existing contracts. This program has been replaced by Software Subscription. But, this Upgrade Protection Plan will not be available for new customer enrollment. Existing contracts for the Upgrade Protection program (5692-ADV) will be honored but will not be renewed, extended, or modified.

14.2 The AIX Bonus Pack - May 7, 2000

AIX Bonus Packs complement the operating system with the benefit of additional packaged software at no additional charge.

An AIX Version 4.3 Bonus Pack is included with every new order when media is selected. In the United States and Canada, updates are available at no

charge for customers who have selected media. An additional media charge for Bonus Pack updates may be required in some countries.

Typical Bonus Pack releases may include:

- Database software
- Development tools
- System management network management utilities
- Multimedia
- Interoperability support
- e-business and commerce
- Country-specific security encryption

Note

Some contents require the latest AIX 4.3 modification level, and contents are subject to change without notice. The contents listed in this document describe the North America edition of the AIX 4.3 Bonus Pack and may not be available in all countries. Encryption subject to import/export restrictions may limit content availability.

Each product is licensed under the terms and conditions that accompany the programs. Installation procedures, terms and conditions, and support may vary. For IBM product software and hardware prerequisites, review the announcement materials located at:

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

14.2.1 Database Software

The following sections describe the available database software located on the AIX Bonus Pack.

14.2.1.1 DB2 Universal Database Version 6.1 for AIX (60-day evaluation software)

This software is the foundation for e-business and is the industry's first multimedia, Web-ready relational database management system. DB2 Universal Database combines power for business intelligence (data warehousing and data mining) with industry-leading performance and reliability to drive the most demanding industry solutions.

Additional product information can be found at:

www.software.ibm.com/data/db2/udb

14.2.1.2 Oracle8i Server

(30 day Trial Software on separate media)

Delivers ease-of-use, power, and excellent performance with a complete suite of point-and-click GUI tools that help speed and simplify database management. Oracle8i Server is currently offered on current RS/6000 Models. Oracle8i Server Enterprise Edition is available on the S-Series family and SP systems.

Additional product information can be found at:

www.oracle.com/database/oracle8i

14.2.2 Development Tools

The following sections discuss the development tools located on the AIX Bonus Pack.

14.2.2.1 Java 2 Technology Edition Version 1.2.2

This product is based on Sun's Java 2 SDK 1.2.2 reference port and includes IBM functional and performance enhancements. The Developer Kit features IBMs Just in Time Compiler (JIT) and Mixed Mode Interpreter as well as a re-engineered Java Virtual Machine (JVM) that passed all tests in the Java Compatibility Kit (JCK).

Additional product information may be found at:

www.ibm.com/developer/java

14.2.2.2 Geodesic Systems Great Circle Version 4.0.6.1

(30 day Evaluation Software)

This is a solution that helps improve performance, availability, and reliability in business-critical server-based systems. It can improve performance and availability in Symmetric Multi-Processor (SMP) environments. Great Circle also includes enhanced components to test and analyze C and C++ applications for memory problems during development and run time.

Additional product information may be found at:

www.geodesic.com/products

14.2.2.3 Java Media Framework Version 1.1.0.1

A cross-platform application programming interface for developing rich media applications, applets, and JavaBeans that runs on all Java-enabled platforms. JMF extends Java's core programming capabilities.

Additional product information can be found at:

java.sun.com/products/java-media/jmf

14.2.2.4 The Kernel Group ZeroFault Dynamic Debugger V2.5 (Evaluation Software with limited error report trace-back)

Uses advanced virtual machine technology to detect and report memory errors and leaks in any C, C++, FORTRAN, Pascal, and Assembly AIX executable, without requiring recompiling, relinking, or access to source code.

Additional product information can be found at:

www.zerofault.com

14.2.2.5 VisualAge for Java, Entry Edition for AIX V2.0.0.3

This VisualAge is a powerful rapid application development tool for building Java-compatible applications, applets and JavaBeans components. The Entry Edition supports a maximum of 500 classes.

Additional product information can be found at:

www-4.ibm.com/software/ad/java

14.2.2.6 AIX Tools Network

Additional tools for application development from IBM and other third party tool developers are available through the AIX Tools Network.

Additional information may be found at:

www.developer.ibm.com/tech/aixtools/index.html

14.2.3 System Management and Network Management Utilities

The following sections discuss the system management and network management utilities included on the AIX Bonus Pack.

14.2.3.1 Tivoli Management Agent (TMA) V3.2

Installation of TMA makes AIX Tivoli Ready. Tivoli Ready products take advantage of key technologies to extend the management services of Tivoli

software to their applications and devices. TMA may be optionally preinstalled.

Additional product information can be found at:

www.tivoli.com/teamtivoli/programs/tivolireadt/agenda.html

14.2.3.2 SecureWay-SSL Version 3.3

This product provides 56-bit encryption support for the IBM SecureWay Directory Version 3.1.1.5 in AIX Version 4.3.3.

Additional product information can be found at:

www.ibm.com/servers/aix/products/aixos

14.2.3.3 SecureWay Directory Server and Client Utilities

(for Maximum Encryption - AIX 4.3.3 only)

Provides 128 bit and Triple DES encryption support for the IBM SecureWay Directory Version 3.1.1 in AIX Version 4.3.3.

14.2.3.4 SecureWay Directory Client Software Developer Kit

(SDK) Encryption Version 3.1.1 (AIX 4.3.3. only)

Provides encryption for 40-bit, 56-bit, and 128-bit encryption support for the IBM SecureWay Directory Version 3.1.1 Software Developer Kit (SDK) supporting Java Naming And Directory Interface (JNDI).

14.2.3.5 IBM IP Security Version 4.3.3

Provides encryption for 40-bit, 56-bit, and Triple DES support for the AIX Version 4.3.3 IP security enhancements.

Additional product information can be found at:

www.ibm.com/servers/aix/products/aixos

14.2.3.6 AIX Certificate and Security Support V4.0

Provides support for the AIX Version 4.3.3 IP security enhancements for the storage of certificates.

14.2.3.7 Web-based System Manager Security V4.3.3

40-bit and 128-bit encryption. Helps provide for the secure operation of the Web-based System Manager servers and clients. It is based on Public Key encryption, the Secure Socket Layer (SSL) protocol, and standard AIX login security.

Additional product information can be found at:

www.ibm.com/servers/aix/products/aixos

14.2.4 Multimedia

The following sections discuss the available multimedia programs included on the AIX Bonus Pack.

14.2.4.1 Ultimedia Services Version 2.3

Ultimedia Services for AIX provides the desktop software tools and utilities for adding audio and video support to AIX. It also provides device control for audio and video adapters on RS/6000 systems.

Additional product information can be found at:

www.ibm.com/servers/aix/products/ibmsw/graphics/ultimedia.html

14.2.4.2 Adobe Acrobat Reader 3.01

The Adobe Acrobat Reader allows you to view, navigate, and print PDF files across all major computing platforms.

Additional product information can be found at:

www.adobe.com/prodindex/acrobat/readstep.html

14.2.5 Interoperability Support

The following interoperability support is included on the AIX Bonus Pack.

14.2.5.1 AIX Fast Connect Version 2.1.1 Evaluation Software

This full-function evaluation copy of AIX Fast Connect for AIX 4.3.3 provides file and print server capabilities to Windows 95, Windows 98, Windows 2000 Windows NT and OS/2 servers and clients. This IBM-developed feature for AIX takes advantage of the scalability and performance capabilities of AIX. An unlimited number of clients are supported.

Additional product information can be found at:

www.ibm.com/servers/aix/products/ibmsw/manage/fastconn.html

14.2.5.2 SCO Tarantella Version 1.4.0.8 Evaluation Software

This product makes applications *Web-enabled* so that they can be published on the Web without rewriting them. It delivers access to any type of application from any Java technology enabled Internet device without requiring any additional software to be installed on that device. This approach

reduces the time to deliver applications and allows the applications to reach more centrally located servers.

Additional product information may be found at:

<http://tarantella.sco.com/>

14.2.5.3 GraphOn GO-Joe Version 2.2.0.0

This product allows Java-enabled devices or browsers to access X applications running on AIX host computers. GO-Joe does not require an X Server to run at the desktop. Instead, GO-Joe's thin server-based architecture permits even the lowest-end Java device to run your X applications as if they were located directly on your desktop.

Additional information can be found at:

www.graphon.com/index.html

14.2.5.4 Novell Network Services (NNS) 4.1 for AIX, Version 2.2.2 (Evaluation Software, 2 users)

An unlimited use license for Novell's latest Novell Directory Services (NDS) for PC clients and browsers. NNS on AIX provides a full-functioned compatible implementation of NDS, Novell distributed file and print services, network security, and administration services across the RS/6000 system family. In addition to its support of Ethernet, token ring, and FDDI LANs, it communicates over IPX/SPX networks and IP networks to support RS/6000 network interoperability with other NetWare and NDS servers. Internet and intranet networks are supported through Novell secure IP tunneling.

Additional product information can be found at:

www.novell.com/press/archive/1997/07/pr97109.html

14.2.6 e-business and Commerce

The following sections discuss the various e-business and commerce solutions available on the AIX Bonus Pack.

14.2.6.1 IBM WebSphere Application Server Version 3.0

This product manages and deploys dynamic Web sites while leveraging open Java-based technologies and application programming interfaces (APIs) as well as the latest XML technologies. This Web application server provides deployment and management of Java, Java Server Pages (JSP) technology and XML applications. The Standard Edition includes the new IBM HTTP

Server Version 1.3.6, (powered by Apache). IBM has enhanced the HTTP server with performance and SSL for secure transactions. When serving static content, the HTTP Server may also see up to 40 percent performance improvement when used with the in-kernel HTTP Get Engine in AIX Version 4.3.3. This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).

Additional product information can be found at:

www.software.ibm.com/websphere/appserv (for application server)

www.software.ibm.com/websphere/httpserver (for http server)

14.2.6.2 Chili!Soft ASP V3.0.3 (Developer's Edition 5 users)

A Web application server based on the popular Active Server Pages (ASP) framework. Chili!Soft ASP enables the development and deployment of ASP applications on AIX and supports ASP applications originally developed for Windows NT.

Additional product information may be found at:

www.chilisoft.com

14.2.6.3 Techexplorer Hypermedia Browser - Introductory Edition (V2.5 Plug-in)

This is packaged with Netscape Communicator and helps extend your Web browser. It can also be turned into an interactive scientific and technical document viewer using TeX or MathML (XML) markup.

Additional product information can be found at:

www.software.ibm.com/enetwork/techexplorer

14.2.6.4 Netscape Communicator 4.7

This product comes with 56-bit and 128-bit encryption with National Language Support. Bi-directional support for viewing Hebrew and Arabic HTML pages is also provided. It also adds the following additional components:

- Netscape Navigator
- Netscape Messenger (Netscape mail)
- Collabra (Netscape discussion groups)
- Composer (Netscape Web page publishing)

Additional product information can be found at:

home.netscape.com

14.2.7 How to Order the Bonus Pack

To receive the AIX Version 4.3 Bonus Pack, contact your IBM Representative, IBM Business Partner, or place an order through Shop IBM. In the United States and Canada, updates are available at no charge for customers who have selected media. An additional media charge for Bonus Pack updates may be required in some countries. When placing an order for the Bonus Pack request 5692-AIX SPO using the information from Table 163, you may need to provide:

- Verification of customer name and number
- Verification of ship-to address
- Number of software shipments to be generated
- RS/6000 hardware and software systems' serial numbers
- Expected shipment/receipt date
- Expedited software delivery requests

Table 163. Bonus Pack Ordering Requirements

Encryption Content	Bonus Pack Feature Codes for 5692-AIX SPO
Includes some products with 40-bit, 56-bit, 128-bit. Triple-DES encryption and certain toolkits for encryption.	0860 US and Canada only
Includes some products with 40-bit, 56-bit, 128-bit. Triple-DES encryption	0905 Worldwide
Includes some products with 40-bit, 56-bit encryption	0861 Worldwide except France
Includes some products with 40-bit, 56-bit encryption	0872 Worldwide

14.3 Third Party Vendors - ISVs

There are over 2000 vendors with applications operating on AIX Version 4.3. For more information, see:

www.software.ibm.com/solutions/isv

Chapter 15. Compilers

This section covers the C, C++, and Fortran compilers and run-time packages available for AIX Version 4. Additional compilers, for a variety of programming languages designed to meet your needs, are available from IBM.

15.1 C Compiler

The levels of C that are currently available from marketing are:

- IBM C for AIX, Version 5 - Supported until January 31, 2002
- IBM C for AIX Version 4.4 - Supported until January 31, 2001
- VisualAge C++ Professional for AIX, Version 5 - Supported until January 31, 2002
- VisualAge C++ Professional for AIX, Version 4 - Supported until January 31, 2001

15.1.1 IBM C for AIX, Version 5 - 5765-E32

Boasting a state-of-the-art C compiler and new innovative graphical source debugger, the IBM C for AIX, Version 5, program is designed to support IBM's powerful AIX Version 4 operating system and latest SMP RS/6000 hardware.

IBM C for AIX, Version 5 extends the existing symmetric multi-processing (SMP) support from IBM C for AIX, Version 4.4, by supporting the OpenMP industry specification. OpenMP provides a model for parallel programming that allows a program to be portable across shared memory architectures from different vendors by using a common set of application program interfaces (APIs). C for AIX provides improved compile and execution performance, additional memory management debug routines, and supports the run-time dynamic linking provided by AIX Version 4.2.1 or later. C for AIX also provides improved prototyping of programs and improved cross-platform compatibility of compilers within the VisualAge family. It offers C language features and improved optimization to protect your investment in existing C applications, including 64-bit integer data type, 128-bit floating-point data type, and run-time address checking. In addition, Version 5 includes a new source debugger with an enhanced graphical user interface (GUI).

IBM C for AIX, Version 5, offers the following features:

- Supports the OpenMP industry specification in 32-bit and 64-bit
- Includes a Distributed Debugger component with an enhanced GUI

- Supports SMP through both automatic and explicit paralleling
- Supports industry standards including ANSI C and Unicode
- Generates highly optimized code for all RS/6000 processors
- Provides profile-directed feedback and interprocedural analysis to optimize performance of C applications
- Provides run-time address checking to detect memory errors
- Adds source code and object code compatibility with previous versions of IBM C for AIX

IBM C for AIX, Version 5, was made available on November 26, 1999. IBM will withdraw this product on January 31, 2002.

15.1.1.1 Hardware and Software Requirements

- Processor: RS/6000 family of processors supported by AIX Version 4.2.1 (5765-C34) or later
- Display: SVGA (800 x 600)
- CD-ROM drive
- Mouse or pointing device
- Memory: 64 MB, or higher, recommended
- Disk space:
 - 80 MB for tools
 - 10 MB for documentation
 - 15 MB for search capability for HTML documents
 - 35 MB for C compiler
 - 64 MB swap/paging space
- IBM AIX Version 4.2.1 (5765-C34), or later, for RS/6000
- X11 required for the Distributed Debugger component
- Frames capable browser to access the HTML-based help, such as Netscape Navigator Version 3.01 or later

IBM C for AIX, Version 5, produces applications that run in the AIX Version 4.2.1 and later environments. Note that 64-bit applications will run only on AIX Version 4.3 and later.

15.1.1.2 Compatibility

C programs written using Versions 3 or 4 of IBM C for AIX are compatible with IBM C for AIX, Version 5.

C programs written using either Version 2 or 3 of IBM C Set ++ for AIX or the XL C compiler component of AIX Version 3.2 are source-compatible with IBM C for AIX, Version 5, with some exceptions to detect unusable programs or areas where results are undefined.

15.1.1.3 Usage Restriction

The IBM C for AIX, Version 5, program is licensed based on authorized users. The number of users allowed to use the C for AIX, Version 5, compiler is determined by the number of licenses purchased by the customer.

An authorized user is a specific user authorized to have local or remote access to the program or any portion of the program. A separate authorization or authorized user is required for each specific user with local or remote access to the program or any portion of the program. Only an authorized user may have access to the program or any portion of the program.

15.1.1.4 Redistribution Information

The files/modules listed below, or located in the directory named below, may be copied onto your media, in object code only, when your application is dependent upon them, subject to the following terms and conditions:

1. Copies of these modules are provided 'AS IS'. You are responsible for all technical assistance for your application.
2. To indemnify IBM from and against any third-party claim arising out of the use or distribution of your application.
3. Not to use IBMs name or trade-marks in connection with the marketing of your applications without IBMs prior written consent.
4. To prohibit the recipient from copying (except for backup purposes), reverse assembling, reverse compiling, or otherwise translating the application.
5. Not to use the same path name as the original files/modules.

An application containing xlsmp.rte or xlsmp.*.rte must be labelled as follows:

"CONTAINS IBM C for AIX, Version 5 Runtime Modules"

15.1.2 IBM C for AIX, Version 4.4 - 5765-C64

IBM C for AIX, Version 4.4, offers C programmers a productive application development environment to construct high-performance, stand-alone, single-machine applications or engineering/scientific applications.

This release is the first IBM C compiler product designed to exploit the RS/6000 SMP architecture. It supports automatic parallelization of a C program as well as explicit parallelization through a set of directives that enables the user to parallelize selected sections of the application program.

IBM C for AIX, Version 4.4, offers the following major advantages:

- Upward source compatibility with previous versions of IBM C for AIX
- An easy, guided migration path from 32-bit to 64-bit applications
- Ability to compile in either 32-bit or 64-bit modes
- SMP support for both 32-bit and 64-bit applications
- Programming support for automatic parallelization
- Programming support for explicit parallelization through directives
- Conformance to industry standards for the C programming language
- Compiler options to enable various levels of optimization for generated object code
- 64-bit-enabled graphical debugger tool

This AIX announcement represents IBM's continuing commitment to the development of applications as demanded by the engineering and scientific markets, with the ability to exploit the power of the RS/6000 architecture.

15.1.3 VisualAge C++ Professional, Version 5 - 5765-E26

The VisualAge C++ Professional for AIX, Version 5, software is a highly productive and powerful development environment for building C and C++ applications. VisualAge C++ Professional offers the flexibility of a set of C++ compilers supported by a suite of tools designed to offer a rich environment for object-oriented application development.

VisualAge C++ Professional for AIX now includes a traditional, makefile-based C++ compiler. This makefile-based compiler supports the latest 1998 ANSI/ISO C++ standard and includes a complete implementation of the ANSI Standard Template Library (STL). The 32-bit and 64-bit architecture and optimization technologies are both supported. RS/6000 symmetric multiprocessor (SMP) architecture is also supported with the

inclusion of the latest IBM AIX C compiler, which supports the C OpenMP industry specification allowing for SMP explicit directives and automatic paralleling.

The VisualAge C++ incremental compiler can give you productivity gains through fast incremental builds and offers you exceptional program-understanding technology to make code development and maintenance easier than ever before.

Additional new features and enhancements in VisualAge C++ Professional for AIX, Version 5, include:

- The new IBM Distributed Debugger for local and remote debugging of 32-bit and 64-bit executables.
- The performance analysis tool to help you understand and improve the performance of C and C++ applications.
- Two new keyboard mappings (vi and emacs, which are popular for AIX) for the editing function within the integrated development environment (IDE).
- Support for multiple code stores within the incremental compiler, allowing a large project to be implemented in separate and smaller units.
- Easier transition from the incremental compiler to the traditional compiler with the ordered name lookup option, which allows the compiler to try to resolve all names to the same declarations that a traditional compiler would.
- More flexible source control with the ability to check files into or out of virtually any source control system.
- Custom interfaces from the IDE to the IBM CMVC and VisualAge TeamConnection source control products.
- Improved runtime performance with the use of a new C++ object mode.
- Better file management support and easier transition from makefiles.

IBM VisualAge C++ Professional for AIX, Version 5 is available at an upgrade price for those customers currently licensed for the following qualifying programs:

- VisualAge C++ Professional for AIX, Version 4 (Product number 5765-D52 or part number 30L8178 (or equivalent part number within the product))
- C and C++ Compilers for OS/2, AIX and for Windows NT, Version 3.6 (Part number 04L3535 (or equivalent part number within the product))

- C Set ++ for AIX, Version 3, product number 5765-421 (All valid feature numbers within the product, except the supply charge or media feature, qualify for the respective upgrade)

Upgrades may be acquired up to the current level of use authorized for the qualifying programs.

15.1.4 VisualAge C++ Professional for AIX, Version 4 - 5765-D52

VisualAge C++ Professional for AIX, Version 4 is a powerful, rapid application development (RAD) tool for building C and C++ applications. The heterogeneous RAD environment provides:

- Tools including a graphical debugger
- Visual Builder
- Data Access Builder
- Incremental compiler and linker
- A rich set of class libraries
- Online help and a powerful full-text search engine

VisualAge C++ Professional provides a standards-compliant C++ compiler. Its incremental development environment and visual programming tools improve programmer productivity. In addition, you can develop applications on AIX and easily port them to OS/2 or Windows NT by leveraging other members of the VisualAge C++ family.

Included with VisualAge C++ Professional is the IBM C and C++ batch compilers for AIX. The IBM C and C++ batch compilers offer 64-bit-enabled batch compilation, memory management and debug memory management routines, a graphical debugger, resource tools for creating and compiling resources, online help, and a powerful full-text search engine.

Table 164 on page 577 shows the C and C++ for AIX products from 1994 onwards. There are only two C and two C++ AIX products available at this time. The others have been withdrawn. Since AIX Version 4, the C Compiler must be ordered separately since it is not included with the operating system.

Table 164 shows the C versions and upgrade paths for the C++ AIX products.

Table 164. C Versions and Upgrade Paths

C /C++ Product	Min. to Max. AIX Level	Upgrade from LPP
5765-E26 VisualAge C++	AIX Version 4.2 or later	5765-D52 VisualAge C++ C and C++ Compilers 3.6 C and C++ Compilers 5765-421
5765-E32 C for AIX, Version 5.0	AIX Version 4.2.1 or later	5765-C64 5765-423 5765-421
5765-D52 VisualAge C++ Professional for AIX Version 4.0	AIX Version 4.1.5 or later	5765-421 C Set ++ for AIX Version 3 5648-A81 C and C++ Compilers for OS/2, AIX and Windows NT Version 3.6
5765-C64 C for AIX Version 4.4	AIX Version 4.1.5 or later	5765-421 C Set ++ for AIX Version 3 5765-423 C for AIX Version 3 5765-C64 C for AIX Version 4.3, Version 4.1
The following products have been withdrawn from marketing:		
5648-A81 C and C++ Compilers for OS/2, AIX, and Windows NT Version 3.6	AIX Version 4.1.4	
5765-C64 C for AIX Version 4.3	AIX Version 4.1.4	5765-421 C Set ++ for AIX Version 3 5765-423 C for AIX Version 3 5765-C64 C for AIX Version 4.1
5765-C64 C for AIX Version 4.1	AIX Version 4.1.4	5765-421 C Set ++ for AIX Version 3 5765-423 C for AIX V3

C /C++ Product	Min. to Max. AIX Level	Upgrade from LPP
5765-421 C Set ++ for AIX Version 3	AIX Version 4.1	5765-035 XL C++ Compiler/6000 Version 1 5765-186 C Set ++ for AIX Version 2 5696-550 C++ POWERbench Version 1 5696-733 C++ POWERbench Version 2 5696-037 SDE Workbench/6000 Version 1 5696-524 SDE Workbench/6000 Version 2 5765-423 C for AIX Version 3
5765-423 C for AIX Version 3	AIX Version 4.1	

15.2 Fortran

This section outlines the various Fortran levels available at AIX Version 4.3 and describes the currently marketed Fortran products, which are:

- IBM XL Fortran for AIX, Version 7.1 (base compiler and base run-time environment)(5765-E02) and XL Fortran Run-Time Environment for AIX, Version 7.1 (base run-time environment only)(5765-E03) - Supported until January 31, 2002
- IBM XL High Performance Fortran Compiler, Version 6.1(5765-612) and XL High Performance Fortran Run-Time Environment, Version 6.1 (5765-613) - Supported until January 31, 2001

15.2.1 IBM XL Fortran for AIX, Version 7.1

IBM XL Fortran for AIX, Version 7.1 (XLF V7.1) is the first XLF compiler to provide support for the full OpenMP Version 1.0 Fortran Application Program Interface. Listed are some features of this product:

- Exploits the RS/6000 symmetric multi-processing (SMP) architecture
- Supports 64-bit pointers and addressability for serial and SMP codes
- Supports the full implementation of the Fortran 95 standard
- Supports automatic parallelization of a Fortran program as well as explicit parallelization (through a set of directives which enable the user to parallelize selected program sections)

- Continues to support both Fortran 90 and FORTRAN 77 programming
- Provides a highly optimized native compiler and a set of AIX-based tools integrated with the AIX Common Desktop Environment for the development of Fortran applications

IBM XL Fortran for AIX, Version 7.1, was made available on March 10, 2000. IBM will withdraw this product on January 31, 2002.

15.2.1.1 Hardware and Software Requirements

The following are the hardware and software requirements for RS/6000 systems and RS/6000 SP systems:

- CD-ROM drive
- 80 MB of RAM
- 70 MB of disk storage
- 64 MB of paging space
- TCP/IP
- AIX Version 4.x, operating system
- AIX windows desktop support
- bos.adt fileset (regardless of the AIX operating system level)
- Frames-capable browser to access the HTML-based online documentation
- Adobe Acrobat reader to view the PDF documentation

Users can also invoke XL Fortran for AIX using an X-station computer or ASCII terminal.

XLF V7.1 includes a technology preview of the IBM Distributed Debugger and continues to ship with the following tools:

- LPEX (a context sensitive editor)
- xldb (a GUI-based debugger)
- xxlf (a GUI-based command line builder)

XLF V7.1 also provides a license management tool that you can customize depending on your business and system environment.

XLF V7.1 is a highly optimized, native Fortran compiler designed to exploit the RS/6000 line of symmetric-multiprocessing (SMP) hardware.

The target market of XLF V7.1 includes the following:

- Scientific and Technical Computing Industries
- Universities
- Research Facilities and Laboratories

XL F V7.1 provides support for:

- OpenMP Version 1.0 Fortran API
- Fortran 90 and 95 standards
- FORTRAN 77 standard

15.2.1.2 Usage Restriction

The license of these products authorizes the customer who already has an IBM XL Fortran for AIX V7.1 program package to make a specific number of additional copies of the product or to allow a specified number of additional users LAN access. Authorizations will be authorized by sending the customer a Proof of License certificate.

The IBM XL Fortran for AIX, Version 7.1, package and upgrade path include one authorized user. The IBM XL Fortran RTE for AIX Version 7.1 package and upgrade path include one authorized install.

15.2.2 XL High Performance Fortran Version 6.1 - 5765-612

IBM XL High Performance Fortran for AIX (XL PF for AIX) is an implementation based on the High Performance Fortran Language Specification, Version 1.1 (High Performance Fortran Forum, November 10, 1994). This implementation provides a standardized set of extensions to Fortran 90 (F90) to improve the performance of a program on parallel computer systems. XL HPF for AIX runs under the AIX, Version 4.1.3, or later, operating system on RS/6000 and SP systems or on a cluster of RS/6000 systems. It exploits the Message Passing Interface Standard (MPI Standard) provided under the IBM Parallel Environment for AIX, Version 2, program to improve the performance of HPF applications.

XL HPF for AIX provides programmers with a highly-optimized native compiler and a set of AIX-based tools integrated with the AIX Common Desktop Environment for the development of Fortran applications. Today's IBM XL HPF for AIX announcement represents IBM's continuing commitment to the development of high-performance applications demanded by the engineering and scientific markets. XL HPF for AIX, Version 1, Release 4, provides selective usability improvements. In addition to the content of XL HPF for AIX, Version 1, Release 3, XL HPF for AIX, Version 1, Release 4, also provides the complete function of XL Fortran for AIX, Version 6 (for

non-HPF codes only), including such new features as symmetric multiprocessor support (automatic or explicit) and support for Fortran 95 (for non-HPF codes only).

IBM XL High Performance Fortran, Version 6.1 (5765-612) and the XL High Performance Fortran Run-Time Environment, Version 6.1 (5765-613) products are available. IBM will withdraw these products on January 31, 2001.

15.2.3 XL Fortran for AIX

Currently, there are two products, XL Fortran RTE for AIX, Version 6.1 (5765-D77) and XL Fortran for AIX, Version 6.1 (5765-D78) in this classification of Fortran.

IBM XL Fortran for AIX, Version 6.1, a follow-on product to XL Fortran for AIX, Version 5.1.1, is the first XL Fortran compiler to provide support for the full Fortran 95 standard. It is also designed to exploit the RS/6000 SMP architecture and to provide support for 64-bit pointers and addressability for serial and SMP codes. XL Fortran for AIX, Version 6.1, provides partial support for the OpenMP specification. Automatic parallelization of a Fortran program as well as explicit parallelization (through a set of directives that enable the user to parallelize selected program sections) is supported. It also continues to support both Fortran 90 and Fortran 77 programming.

XL Fortran for AIX, Version 6.1, provides a highly-optimized native compiler and a set of AIX-based tools integrated with the AIX Common Desktop Environment for the development of Fortran applications. XL Fortran for AIX, Version 6.1, continues to ship with the following tools:

- LPEX (a context sensitive editor)
- xldb (a GUI-based debugger)
- xxlf (a GUI-based command line builder)

XL Fortran for AIX, Version 6.1, also provides a license management tool that you can customize depending on your business and system environment.

Table 165 outlines the Fortran Versions that are orderable today and the AIX levels that they can be installed on.

Table 165. Fortran Versions and Their AIX Levels

Fortran Version	Min. to Max. AIX Level
5765-D77 XL Fortran RTE for AIX, Version 6.1	AIX Version 4.1 through AIX Version 4.3.2

Fortran Version	Min. to Max. AIX Level
5765-D78 XL Fortran for AIX, Version 6.1	AIX Version 4.1 through AIX Version 4.3.2
5765-E02 XL Fortran for AIX, Version 7.1	AIX Version 4.3.3 (works on 4.x)
5765-E03 XL Fortran RTE for AIX, Version 7.1	AIX Version 4.3.3 (works on 4.x)
The following products have been withdrawn from marketing:	
5765-612 XL High Performance Run-Time Version 1.3, Version 1.4	AIX Version 4.1 through AIX Version 4.3.2
5765-613 XL High Performance Fortran Version 1.3, Version 1.4	AIX Version 4.1 through AIX Version 4.3.3
XL Fortran for AIX, Version 5.1.1	AIX Version 4.1 through AIX Version 4.3.2
XL Fortran RTE for AIX, Version 5.1.1	AIX Version 4.1 through AIX Version 4.3.2
5696-551 Fortran POWERbench 1.0	AIX Version 4.1 or later
5765-018 Fortran Compiler	AIX Version 4.1 or later
5765-019 Fortran Run-Time Environment	AIX Version 4.1 or later
5765-526 XL Fortran RTE Version 3	AIX Version 4.1 or later
5765-657 XL Fortran RTE for AIX, Version 4	AIX Version 4.1 or later
5765-658 XL Fortran for AIX, Version 4	AIX Version 4.1 or later
5765-612 XL High Performance Fortran Version 1	AIX Version 4.1 or later
5765-613 XL High Performance Fortran Version 1	AIX Version 4.1 or later

Chapter 16. High Availability

Until recently, the only method for achieving high availability in the UNIX realm was through fault tolerant technology. Fault tolerance relies on specialized hardware to detect a hardware fault and instantaneously switch to a redundant hardware component, whether the failed component is a processor, memory board, power supply, I/O subsystem, or storage subsystem. Although this cutover is apparently seamless and offers non-stop service, a high premium is paid in both hardware cost and performance because the redundant components do no processing. More importantly, the fault tolerant model does not address software failures, which are the most common reason for down time.

High availability (HA) views availability, not as a series of replicated physical components, but rather as a set of system-wide, shared resources that cooperate to provide essential services. High availability combines software with industry-standard hardware to minimize down time by quickly restoring essential services when a system, component, or application fails. While not instantaneous, services are restored rapidly, often in less than a minute.

Figure 92 shows the costs of downtime per minute from a recent study.




Figure 92. Costs of Downtime

The availability of a system is dependent on a wide range of factors. Some of the key factors that can significantly impact system availability include the following:

- Hardware Reliability
- Software Reliability
- System Management Methods and Operational Standards
- System Administrator Skills
- Data Storage Management
- Environmental Factors

The following section provides an overview of the high availability products for the AIX and RS/6000 family:

- High Availability Cluster Multi-Processing (HACMP)
- High Availability Geographic Cluster/Geographic Remote Mirroring (HAGEO/GeoRM)

16.1 High Availability Cluster Multiprocessing

HACMP was originally developed to enhance the reliability of AIX and RS/6000 hardware to meet the demands of mission-critical workloads.

HACMP aims, in a well designed cluster, to eliminate any single point of failure, be it disks, network adapters, power supply, processors, or other components.

Note

Cluster multi-processing is a group of loosely coupled machines networked together sharing disk resources. In a cluster, multiple server machines cooperate to provide a set of services or resources to clients.

An HACMP cluster is made up of the following components:

- Nodes
Nodes form the core of an HACMP cluster. A node is a processor that runs both AIX and the HACMP software. The HACMP software supports RS/6000 uniprocessor and symmetric multiprocessor (SMP) systems and the SP nodes as cluster nodes.
- Shared external disk devices

Each node must have access to one or more shared external disk devices. A shared external disk device is a disk physically connected to multiple nodes. The shared disk stores mission-critical data, typically mirrored for data redundancy. A node in an HACMP cluster must also have internal disks that store the operating system and application binaries, but these are not shared.

- Networks

Nodes in an HACMP cluster use the network to allow clients to access the cluster nodes, enable cluster nodes to exchange keep-alive messages, and in concurrent access environments, serialize access to data.

Types of networks:

Public network Connects multiple (two or more) nodes and allows clients to access the cluster nodes. Ethernet, token ring, and FDDI networks can be defined as public networks. A SLIP line, which does not provide any client access, can also be defined as a public network.

Private network Provides point-to-point communication between two nodes. It typically does not allow client access. HACMP uses a public network for lock traffic if no private networks are available. Ethernet, token ring, FDDI, Serial Optical Channel Connector (SOCC), and Asynchronous Transfer Mode (ATM) networks can be defined as private networks.

Serial network Provides a point-to-point, non-IP connection between two cluster nodes for HACMP for AIX control messages and heartbeat traffic in the event the TCP/IP subsystem fails.

- Network adapters

Typically, a node should have at least two network adapters (a service adapter and a standby adapter) for each connected network.

- Clients

A client is a processor that can access the nodes in a cluster over a public local area network. Clients each run a *front end* or client application that queries the server application running on the cluster node.

Figure 93 on page 586 shows a typical HACMP solution.

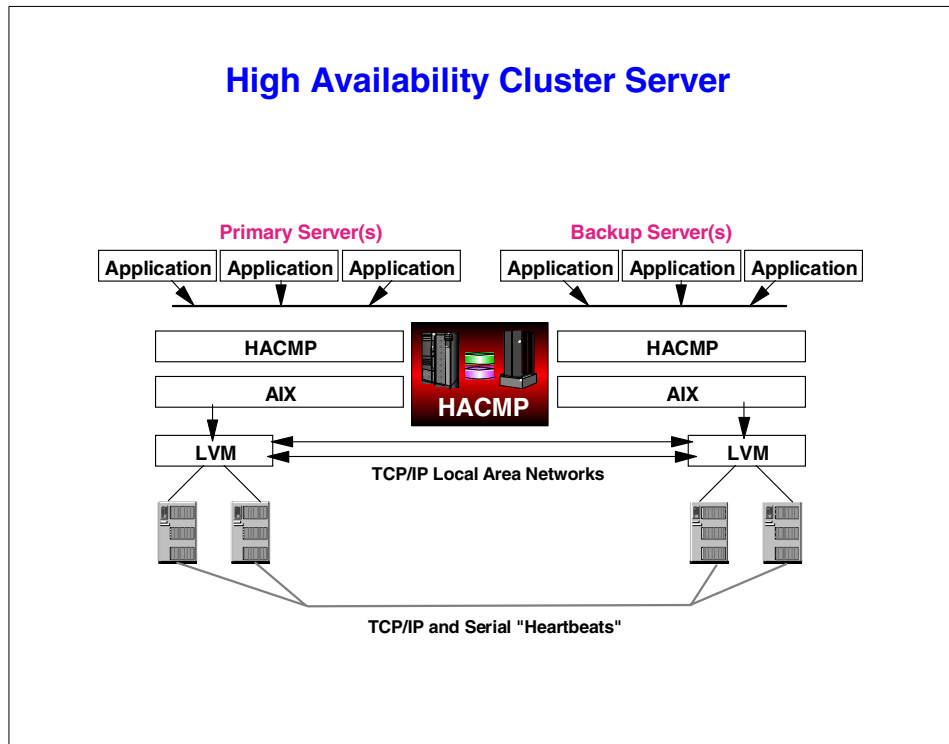


Figure 93. Typical HACMP Solution

HACMP provides a highly-available environment by:

- Identifying the set of cluster resources that are essential to processing.
- Defining the takeover relationships among the cluster nodes that access these resources.

By identifying resources and defining takeover relationships, HACMP makes numerous cluster configurations possible, providing tremendous flexibility in defining a cluster environment tailored to individual requirements.

Detecting the failure of a system or an adapter is achieved by sending *heart beats* (small messages that indicate the sender is functioning correctly) between systems. Loss of heart beats for an extended period of time can be reasonably interpreted as meaning that the sending system has failed.

Cluster resources can include both hardware and software, such as:

- Disks

- Volume groups
- File systems
- Network addresses
- Application servers

To be made highly-available by HACMP, each resource must be included in a resource group. Resource groups allow related resources to be combined into a single logical entity for easier configuration and management.

The takeover relationships among cluster nodes determine which cluster nodes control a resource group and which cluster nodes take over control of the resource group when the original node relinquishes control. Define the takeover relationship of a resource group by assigning it one of the following type designations:

- Cascading

A cascading resource group defines a list of all the nodes that can control the resource group and then, by assigning a takeover priority to each node, specifies a preference for which cluster node controls the resource group. When a failover occurs, the active node with the highest priority acquires the resource group. If that node is unavailable, the node with the next-highest priority acquires the resource group, and so on.

- Rotating

A rotating resource group, such as a cascading resource group, defines the list of nodes that can take over control of a resource group and uses priorities to determine the order in which other nodes can take control of the resource.

- Concurrent

A concurrent access resource group may be shared simultaneously by multiple nodes. All nodes concurrently accessing a resource group acquire that resource group when they join the cluster. There are no priorities among nodes.

Additional information how high availability is achieved and what is possible in a cluster can be found in *Concepts and Facilities Guide*, shipped with the HACMP and HACMP ES products.

An important issue for setting up an HA environment is the planning phase. Considerable work has been done to make the planning process easier and to try and eliminate sources of error. For additional information, read *HACMP V4.3: Planning Guide*, SC23-4277.

Management and maintenance of the cluster, once it is up and running, is another aspect of the HACMP software that has been significantly developed to make it easier to set up and manage a complex HA environment. The *HACMP Administration Guide* documents the common administration tasks and tools that exist to help carry out these tasks.

Among the tools that comprise HACMP and HACMP ES, the following probably contribute most to the manageability of a multi-node environment:

- Version compatibility to support node-by-node migration for upgrade or migration of HA software.
- Monitoring tools, clstat for cluster status, event error notification, the add on HAView (for use with the Netview product), and log files.
- Cluster single point of control (C-SPOC), which allows an operation defined on one node to be executed on all nodes in the cluster (to maintain shared logical manager components, users, and groups).
- Dynamic reconfiguration of resources and cluster topology (DARE).
- The verification tool, clverify, which can be run after any changes are made to the cluster to verify that the cluster topology and resource settings are valid.
- The cluster snapshot utility that can be used to save and restore cluster configurations.

16.1.1 Features in HACMP

HACMP has been in the market place since 1992 when Version 1.1 of the high availability product HACMP/6000 was announced. This was followed by Version 1.2, Version 2 and Version 2.1, and Version 3, Version 3.1, and Version 3.1.1, all of which were based on AIX 3.2.5. Since this time, there have been three releases of HACMP Version 4 (4.1, 4.2, 4.3, and 4.4). These are referred to as HACMP *Classic* which can support up to eight nodes in a cluster.

The product line has been expanded to include HACMP Enhanced Scalability (HACMP ES), which can support up to 32 nodes in a cluster.

The brief history of HACMP versions can be found in Appendix A, "HACMP History" on page 625.

Scalability, support of large clusters, and, therefore, large configurations of nodes (and potentially disks), leads to a requirement to manage clusters of nodes. To address management issues and take advantage of new disk attachment technologies, HACMP ES was released.

This was originally only available for the SP, where tools were already in place with PSSP to manage larger clusters.

The benefits of HACMP ES arise from the technologies that underlie its heart beating mechanism supplied in RISC System Cluster Technology (RSCT). HACMP ES can be used to create a large HA cluster and take advantage of, for example, disk technologies, such as storage area networks (SANs), that allow disks to be accessed by a large number of systems without a performance penalty.

The subsections that follow provide a brief look of HACMP and HACMP ES. Before this, however, it is useful to know the features that comprise HACMP and HACMP ES:

- HACMP (Classic)
 - High availability subsystem (HAS)
 - Concurrent resource manager (CRM)
 - High availability network file system (HANFS)
- HACMP ES
 - Enhanced scalability (ES)
 - Enhanced scalability concurrent resource manager (ESCRM)

16.1.2 HACMP and HACMP ES Version 4.3

Custom snapshot methods were introduced in HACMP 4.3.0 in September, 1998, allowing an administrator to capture additional information with a snapshot. The ability to capture a cluster snapshot was added to the HACMP ES feature.

- C-SPOC support for shared volumes and users or groups was further extended. All nodes in a cluster are now immediately informed of any changes; so, failover time is reduced, as “lazy update” will not usually need to be used.
- A task guide was introduced for shared volume groups that reduces the time spent on some administration tasks within the cluster (`smit hacmp -> Cluster System Management -> Task Guide For Creating A Shared Volume Group`).
- Multiple pre- and post-event scripts are now supported in 4.3.0, therefore, extending once more the level to which an HACMP or HACMP ES cluster event can be tailored.

- HACMP ES is now available on stand-alone RS/6000s. RISC system cluster technology (RSCT), previously packaged with PSSP, is now packaged with HACMP ES.
- Scalability in HACMP ES was increased once more to support up to 32 nodes in a SP cluster. HACMP *Classic* was still limited to eight nodes.
- The CRM feature was added to HACMP ES. Eight nodes in a cluster are supported for ESCRM.

16.1.3 HACMP and HACMP ES Version 4.3.1

HACMP 4.3.1 and HACMP ES offers improved usability, more flexible installation options, and additional hardware and software support for RS/6000 customers with mission-critical applications.

HACMP 4.3.1 has five major components:

- High Availability Subsystem (HAS)

HAS provides the base services for cluster membership, system management, configuration integrity, and control, failover, and recovery. Cluster status and monitoring facilities are also there for the programmer and system administrator. The High Availability Subsystem contains all of the features described in the remainder of this section, with the exception of the Concurrent Resource Manager, HANFS for AIX, Enhanced Scalability, and Enhanced Scalability Concurrent Resource Manager, which are separately ordered components.
- Concurrent Resource Manager (CRM)

CRM optionally adds concurrent shared-access management for supported RAID and SSA disk subsystems. Concurrent access is provided at the raw disk level. The application must support some mechanism to control access to the shared data, such as locking. HAS provides distributed locking facilities to support this. The CRM component includes the High Availability Subsystem component.
- High-Availability for Network File System for AIX (HANFS for AIX)

HANFS for AIX optionally provides high availability for data accessed via the Network File System (NFS). HANFS for AIX is intended to replace the HANFS Version 3 feature for customers who are using AIX Version 4.1.5, or later. This component does not include the High Availability Subsystem, the Concurrent Resource Manager, or the Enhanced Scalability components, nor can it be executed in the same cluster as these components or installed on a system containing these components.
- Enhanced Scalability (ES)

ES builds on the Event Management and Group Services facilities of Parallel Systems Support Programs (PSSP) for AIX to scale HACMP up to 32 SP nodes. It also allows customers to use these facilities to define their own HACMP events. The Event Management and Group Services facilities provide detection and notification of the loss of a node, network, or adapter. ES then drives the appropriate HACMP recovery action. The customer can similarly define a recovery program and have ES execute it in response to any event that Event Management can process using the resource monitors supplied with PSSP Version 2.3. ES does not include the High Availability Subsystem, Concurrent Resource Manager, or HANFS for AIX features.

- **Enhanced Scalability Concurrent Resource Manager (ESCRM)**

ESCRM optionally adds concurrent shared-access management for supported RAID and SSA disk subsystems. Concurrent access is provided at the raw disk level. The application must support some mechanism to control access to the shared data, such as locking. The ESCRM component includes the ES component and the HACMP distributed lock manager.

16.1.3.1 HACMP and HACMP ES 4.3.1. Features

Support for target mode SSA was added to the HACMP ES and ESCRM components with version 4.3.1 in July, 1999.

- HACMP 4.3.1 supports the AIX Fast Connect application as a highly available resource, and the `clverify` utility is enhanced to detect errors in the configuration of AIX Fast Connect with HACMP.
- The dynamic reconfiguration resource migration utility is now accessible through SMIT panels as well as the original command line interface, therefore, making it easier to use.
- There is a new option in SMIT to allow, in software, an active service or boot adapter to swap to a standby adapter (`smit hacmp -> Cluster System Management -> Swap Adapter`).
- The location of HACMP log files can now be defined by the user by specifying an alternative path name via SMIT (`smit hacmp -> Cluster System Management -> Cluster Log Management`).
- Emulation of events for AIX error notification.
- Node-by-node Migration to HACMP ES is now included to assist customers who are planning to migrate from the HAS or CRM features to the ES or ESCRM features, respectively. Prior to this, conversion from HACMP to HACMP ES was the only option and required the cluster to be stopped.

- C-SPOC is again enhanced to take advantage of improvements to the LVM with the ability to make a disk known to all nodes and to create new shared or concurrent volume groups.
- An online cluster planning work sheet program is now available.
- The limitations detailed for HACMP ES 4.2.1 and 4.2.2 have been addressed with the exception of the following:
 - Visual System Manager (VSM) is not supported with a cluster greater than eight nodes.
 - Enhanced Scalability cannot execute in the same cluster as the HAS, CRM, or HANFS for AIX features.
 - The forced down function is not supported.
 - SOCC and SLIP protocols are not supported

Note

The HACMP cluster configuration is determined by the number of nodes and the way the resources are defined. Up to eight nodes are supported with the HAS, CRM, and ESCRM features, and up to 32 nodes are supported with the Enhanced Scalability feature.

With clusters of up to 32 nodes, the inherent configuration flexibility is tremendous:

- SCSI disks can be attached to up to four nodes
- SSA disks can be attached to up to eight nodes

Note

The HACMP ES 4.3 manuals are available at the Web page:

http://www.rs6000.ibm.com/doc_link/en_US/a_doc_lib/aixgen/hacmp_index.ht

16.1.4 HACMP Version 4.4

HACMP Version 4.4 was introduced June 23, 2000 for RS/6000 uniprocessor and SMP servers and July 28, 2000 for RS/6000 SP systems. It provides the following enhancements to the previous release:

- Application Monitoring (ES and ESCRM only) provides Process Application monitoring and User-defined Application monitoring to determine the state of an application and to restart the application or fall the resource group over to another node.

- Tivoli Cluster Monitoring allows users to monitor the state of an HACMP cluster and its components on a Tivoli Desktop window.
- Cascading without Fallback permits specifying that the resource group not return to the original node when that node rejoins the cluster.
- Cluster Verification is enhanced to detect additional startup of fallover problems.
- New Performance Tuning Parameters provide easier and more granular control.
- New documentation provides guidance for setting up, monitoring, and managing 24 x 7 clusters.
- Enhanced LVM Task Guide now provides a display of the physical location of each available disk and will create automatically a JFS log file.
- NFS Migration (HAS only) provides for migration between HACMP 4.3.1 HANFS feature and HACMP 4.4 HAS feature. HANFS is no longer included as a feature of HACMP.
- CSPOC is enhanced for specifying a user password and file system creation.

16.1.5 Upgrade and Migration Paths

The need to keep software at reasonably current levels arises partly from the fact that software products are continually being developed to include new features, improve existing features, and also to take advantage of new technology.

16.1.5.1 Version Compatibility

Version compatibility was introduced with HACMP 4.2 to allow node-by-node migration to a later release of HACMP. During the migration, the cluster will comprise nodes at two different levels of HACMP. Version compatibility supports this cluster state. The rest of the cluster nodes are migrated, one at a time, until all nodes in the cluster are running the same level of HACMP. The concept of version compatibility also applies in an HACMP ES environment.

16.1.5.2 Classic-to-Classic

Table 166 shows, for each release of HACMP Classic, whether fixes or install media are required to upgrade or migrate to subsequent levels of HACMP.

Note

Each of the following options are referred to as HACMP Classic:

- HAS
- CRM
- HANFS

Table 166. Upgrade Path or Migration for HACMP Classic

From/To	4.1.1	4.2.0	4.2.1	4.2.2	4.3.0	4.3.1	4.4
4.1.0	M	M	M	M	M	M	M
4.1.1	/	M	M	M	M	M	M
4.2.0	/	/	P/M	P/M	M	M	M
4.2.1	/	/	/	P/M	M	M	M
4.2.2	/	/	/	/	M	M	M
4.3.0	/	/	/	/	/	P*/M	M
4.4	/	/	/	/	/	/	M

M can migrate to the given level of HACMP via install media.
P can upgrade to the given level of HACMP via PTF (fix package).
* PTF U465440

It should also be noted that some of the devices supported by a given level of HACMP may require a later release of the operating system. Table 167 shows the compatibility of HACMP with AIX Version 4.

Table 167. Compatibility of HACMP with Supported AIX Version 4

HACMP Classic	Level of AIX
4.2.1	4.2.1, 4.2.2
4.2.2	4.2.1, 4.2.2, or 4.3
4.3.0	4.3.2, 4.3.3
4.3.1	4.3.2, 4.3.3
4.4	4.3.3

16.1.5.3 HACMP ES to HACMP ES

Table 168 shows, for HACMP ES, whether fixes (PTFs) or install media is required to move from a given release of HACMP ES to another.

Note

Each of the following options is referred to as HACMP ES:

- Enhanced Scalability (ES)
- Enhanced Scalability Concurrent Resource Manager (ESCRM)

Table 168. Upgrade or Migration for HACMP ES

From/To	4.2.2	4.3.0	4.3.1
4.2.1	P/M PSSP 2.3, or later	M PSSP 3.1, or later	M PSSP 3.1.1, or later
4.2.2	/	M PSSP 3.1, or later	M PSSP 3.1.1, or later
4.3.0	/	/	M* PSSP 3.1.1, or later

M can migrate to the given level of HACMP ES via install media.
P can upgrade to the given level of HACMP ES via PTF (fix package).
* For upgrade from 4.3.0 to 4.3.1, HACMP ES needs to order a media refresh.
NOTE: PSSP for SP environments only.

Where HACMP ES is installed on the SP, there is the added consideration of compatibility with the level of Parallel Systems Support Programs (PSSP) installed. The level of PSSP required to support a given version of HACMP ES can have implications for upgrades beyond the nodes in the HA cluster.

If the level of PSSP required is higher than that currently installed on the control workstation, it must be upgraded first; however, it is not necessary to update all nodes in the SP frame, only those that are to be part of the HA cluster. Issues with PSSP must be resolved before upgrading or migrating HACMP ES on the SP.

Table 169 shows the compatibility of HACMP ES with AIX operating systems.

Table 169. Compatibility of HACMP ES with AIX Version 4

HACMP ES	RS/6000 SP	RS/6000	Level of AIX
4.2.1	Yes	No	4.2.1
4.2.2	Yes	No	4.2.1 or 4.3.x
4.3.0	Yes	Yes	4.3.2, 4.3.3
4.3.1	Yes	Yes	4.3.2, 4.3.3

HACMP ES	RS/6000 SP	RS/6000	Level of AIX
4.4	Yes	Yes	4.3.3

HACMP ES 4.2.2, and earlier, is not supported on stand-alone RS/6000s.

16.1.5.4 Classic to ES

When migrating from HACMP Classic to HACMP ES, there is a difference depending on whether the system is an SP or a stand-alone RS/6000.

Table 170 shows the possible migration from HACMP Classic to HACMP ES.

Table 170. Migration from HACMP Classic to HACMP ES

From Classic To ES	4.2.1	4.2.2	4.3.0	4.3.1	4.4
4.1.0	*(SP only)	*(SP only)	M	M	M
4.1.1	*(SP only)	*(SP only)	M	M	M
4.2.0	*(SP only)	*(SP only)	M	M	M
4.2.1	M (SP only)	*(SP only)	M	M	M
4.2.2	/	M (SP only)	M	M	M
4.3.0	/	/	M	M	M
4.3.1	/	/	/	M	M
4.4	/	/	/	/	M

M can migrate to the given level of HACMP ES via install media.
 SP only is reported where the level of HACMP ES is only available for the SP system.
 *Must first upgrade to a comparable level of HACMP, then conversion to HACMP ES is supported via install media.

There are two methods that can be used to migrate from HACMP Classic to HACMP ES listed below and described in more detail in *Migrating to HACMP/ES*, SG24-5526:

- Node-by-node migration from HACMP Classic for AIX to HACMP ES, 4.3.1, or later.
- Convert from HACMP Classic to HACMP ES using a snapshot of the Classic cluster.

16.1.5.5 Unsupported Levels of HACMP

The earlier versions of HACMP, Versions 1.1, 1.2, 2.1, 3.1, and 3.1.1, are no longer supported, and there is no direct migration path to the current level of HACMP. These earlier versions of the code were based on AIX, Version 3.2.5. HACMP Version 4.1, the first HACMP release to be based on AIX, Version 4, is no longer under defect support, but a migration path exists to the latest level of HACMP and HACMP ES. If there is a cluster at an unsupported level of HACMP with no migration path, then it must be un-installed prior to installing and configuring a later release. This implies thorough planning as per the *HACMP AIX Planning Guide* for HACMP or HACMP ES. Device support is backwards compatible; so, the devices that comprise the cluster will be supported under a more current version of HACMP or HACMP ES.

16.1.6 High Availability Geographic Cluster Product

High Availability Geographic Cluster (HAGEO) provides a functional extension to HACMP that enables the rapid recovery of system and data resources in the event of a major disaster, such as an extended power outage, fire, flood, hurricane, or earthquake. Designed to operate over unlimited distances, HAGEO substantially reduces business risks associated with an outage due to a computer room, building, or site-wide catastrophe. HAGEO and HACMP together allow a customer to protect their data from both the failure from any single component within a local HACMP cluster and also from the failure of a single site.

HAGEO was announced in 1996. In 1999, Geographic Remote Mirroring (GeoRM) was announced as a new offering with all of the data mirroring function of HAGEO but without the automated failure detection, recovery, and reintegration of HAGEO.

The following sections cover these products in more detail.

Figure 94 on page 598 shows an example of an HAGEO solution.

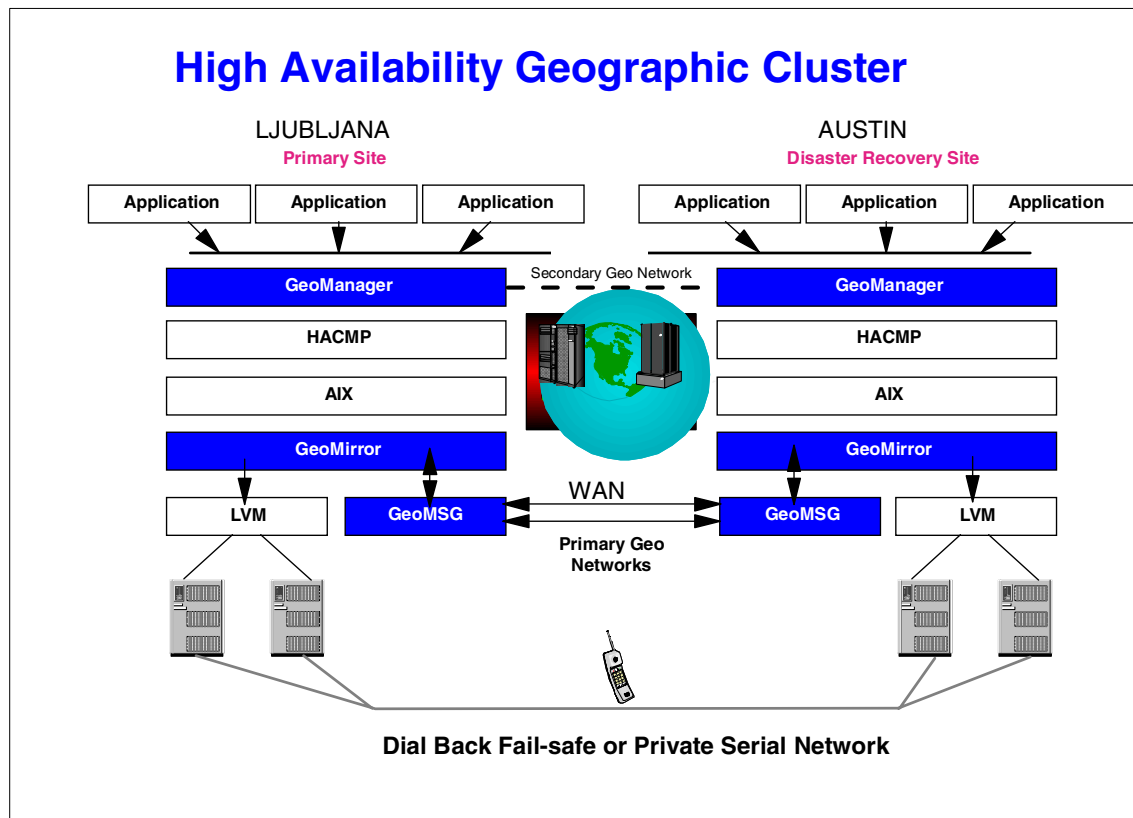


Figure 94. HAGEO Solution

16.1.6.1 HAGEO Version 2.1

IBM increases the number of applications running on RS/6000 with the addition of IBM High Availability Geographic Cluster for AIX, Version 2.1 (HAGEO 2.1). HAGEO 2.1 provides real-time mirroring of customer data between systems connected by local or point-to-point networks, bringing disaster recovery capability to an RS/6000 cluster placed in two widely separated geographic locations.

HAGEO 2.1 allows a cluster of RS/6000 computers to each maintain an exact replica of the application and data. Each site in this geographic cluster can be a single system or multiple systems, with the entire geographic cluster containing no more than eight nodes. HACMP provides automatic detection and failover between sites. HACMP also provides the protection within a site for HAGEO 2.1 by reacting to that sites component failures and by preventing unnecessary geographic takeovers.

HAGEO 2.1 has three functional components:

- Geographic Cluster Manager (GeoManager)
- Geographic Mirroring (GeoMirror)
- Geographic Messaging (GeoMessage)

GeoManager acts as a supervisor for the distributed cluster. It reacts to network and site related events, such as a site failure or a site isolation.

GeoMirror works on top of the AIX logical volume disk manager. It replicates data and coordinates updates at each location, ensuring that sites are synchronized.

GeoMessage provides reliable communication between sites.

HAGEO 2.1 operates in one of three possible modes:

- Synchronous
- Synchronous with Mirror Write Consistency (MWC)
- Asynchronous

In synchronous mode, data is written first to the remote site and then at the local site where the transaction was originally initiated. This mode offers the mirror image with the least amount of time between the writes at each site. Synchronous mode ensures that the same data exists on the disks at both sites at the completion of every write.

In synchronous with MWC mode, the GeoMirror device operates as it does in synchronous mode, but also keeps a detailed record in its state map of all requests that have not completed at both local and remote sites. This guarantees that both sites can be restored to have identical copies of the data, even in the event of a site failure in mid-transaction.

In asynchronous mode, the GeoMirror device does not wait for the write to the remote disk to complete before it returns to the application. The writes to the remote disk, therefore, lag behind the writes to the local disk. This mode offers the best performance on the local machine, however, a resulting longer period of time exists when the data on the local disk and the remote mirror differ. In the event of a site failure, some transactions may be lost.

Ultimately, customers should decide what mode is appropriate for their HAGEO cluster based on the trade-off between exact mirroring of data versus application performance.

HAGEO 2.1 provides a data and system resource replication environment that extends HACMP clustering function to a larger area by enabling mirroring, failover, cluster heartbeat, and event management over a variety of LAN and point-to-point communication adapters, including token ring, Ethernet, FDDI, T1, and ATM.

Specifically designed for fast recovery, HAGEO 2.1 differs from less effective disaster protection schemes by recovering application, data, and system environments in minutes rather than in hours or even days. The components of HAGEO 2.1 offer the following features:

- Mirroring for storage over local area and point-to-point networks
- Synchronous or asynchronous data backup
- Protection of data and applications against total site failure

16.1.6.2 Geographic Remote Mirroring (GeoRM) Version 2.1

For customers who require only remote data mirroring, GeoRM provides a potentially less complex and lower cost mirroring environment than that provided by existing IBM HACMP/HAGEO offerings. If automated failure detection and recovery is required, the HAGEO offering should be proposed in preference to GeoRM. GeoRM includes the GeoMirror and GeoMessage functional components of HAGEO.

Access to mirrored data, when using GeoRM following a failure, will require manual intervention. Customers who desire to partially automate this process, for example, by using shell scripts, should either have sufficient skills to do this themselves or be prepared to purchase services to have these shell scripts and programs created for them.

GeoRM is positioned at the cost-effective end of the disaster recovery product spectrum. GeoRM can be configured to mirror the data to be protected, regardless of the application and location, and can provide complete remote mirroring of data.

GeoRM 2.1 can mirror data from up to seven RS/6000 or RS/6000 SP-based servers to a single remote server, with no limitation on the distance between locations. Since GeoRM 2.1 is both database and file-system independent, it is not necessary to modify applications that utilize GeoRM's mirroring capabilities. A server at a geographically remote location can mirror customer data in real time using a LAN or point-to-point networks. GeoRM 2.1 is designed to prevent a total failure at one location from causing a loss of data at the remote location.

The announcement of GeoRM provides customers with an alternate data recovery strategy with the following features:

- A cost-effective data mirroring solution providing remote backup of customer-critical business data for applications running across the RS/6000 product family
- Mirroring for storage over local area and point-to-point networks
- Protection of data against total location failure
- *Many to one* mirroring
- Synchronous, asynchronous, and Synchronous with Mirror Write Consistency (MWC) data backup modes.

The following Web site contains useful information on high availability:

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

16.1.7 Summary

Table 171, and Table 172 on page 603, list supported systems, devices, or features with HACMP versions.

Table 171. Supported Systems, Devices, or Features with HACMP Versions

	HACMP 4.2.2 AIX 4.2.1	HACMP 4.2.2 AIX 4.3.2	HACMP 4.2.2 AIX 4.3.3
VSS 2105-B09 2105-100	APARs IX81408-HA 1X62304-AIX	APAR IX81408-HA	APARs IX81852-FCAL IX818520FCSS
SSA 7133-D40 7133-T40	Yes	Yes	Yes
43P 7043-150 7043-260	Yes	Yes	Yes
SP Node 9076-50H 9076-55H	No	PSSP 3.1	PSSP 3.1
Gigabit Ethernet (# 2969)	No	Yes	Yes
8-Port Async (# 2943)	Yes	Yes	Yes
128-Port Async (# 2944)	Yes	Yes	Yes
10/100 Mb Ethernet	Yes	Yes	Yes
Advanced SSA/RAID PCI Adapter (# 6225)	Yes	Yes	Yes
ESS 2105-E10 2105-E20	APAR: IY04403 DPO Not Supported CRM Requires APAR: U458552 Limit 4 nodes/cluster	APAR: IY04403 DPO w/CRM only CRM Requires APAR: U458552 Limit 4 nodes/cluster	APAR: IY04403 DPO w/CRM only CRM Requires APAR: U458552 Limit 4 nodes/cluster
7017-S80	No	No	Yes
7017-S70	No	Yes	Yes
4/16 token ring (# 4959)	Yes	Yes	Yes
Dual Ultra2 SCSI PCI (# 6205)	Yes	Yes	Yes
2104-DL1/TL1	Yes	Yes	Yes
Gigabit Fibre Channel ¹ (# 6227)	No	No	Yes
2102-F10 ¹ 2102-D00 ¹	No	No	Yes

	HACMP 4.2.2 AIX 4.2.1	HACMP 4.2.2 AIX 4.3.2	HACMP 4.2.2 AIX 4.3.3
Advanced SerialRAID Plus (# 6230)	Yes	Yes	Yes
32 MB Fast Write Cache (# 6235) (2-initiator only)	Yes	Yes	Yes
7044-170	No	No	Yes
7044-270	No	No	Yes
Ultra SCSI F/W # 6204	No	No	Yes
4-Port 10/100 Ethernet (# 4951)	No	No	Yes
375 MHz POWER3 Node	No	No	PSSP 3.1.1 APAR IY06844
3 Channel UltraSCSI RAID (#2494)	No	No	No
2 Port Multi Protocol PCI adapter (# 2962)	No	No	No
ESS 2105-F10 2105-F20	No	No	No
2108-G7 SAN Data Gateway	No	No	No
¹ Not supported in HA SP environments			

Table 172. Supported Systems, Devices, or Features with HACMP Versions

	HACMP 4.3 AIX 4.3.2	HACMP 4.3.1 AIX 4.3.3	HACMP 4.4 AIX 4.3.3
SSA 7133-D40 7133-T40	Yes	Yes	Yes
VSS 2105-B09 2105-100	Yes	APARs IX81852-FCAL IX81852-FCSS	Yes
43P 7043-150 7043-260	Yes	Yes	Yes
SP Node 9076-50H 9076-55H	PSSP 3.1	PSSP 3.1	PSSP 3.2
Gigabit Ethernet (# 2969)	Yes	Yes	Yes
8-Port Async (# 2943)	Yes	Yes	Yes

	HACMP 4.3 AIX 4.3.2	HACMP 4.3.1 AIX 4.3.3	HACMP 4.4 AIX 4.3.3
128-Port Async (# 2944)	Yes	Yes	Yes
10/100 Mb Ethernet	Yes	Yes	Yes
Advanced SSA/RAID PCI Adapter (# 6225)	Yes	Yes	Yes
ESS 2105-E10 2105-E20	No	APARs IY03438, IY11560, IY11564 DPO w/CRM only CRM requires APAR: U458552 Limit 4 nodes/cluster	APARs IY11563, IY11565
ESS 2105-F10 2105-F20	No	HACMP APARs IY11110, IY11560, IY11564	Yes, APARs IY11480, IY11563, IY11565
2108-G7 SAN Data Gateway	No	APAR IY07313 APAR IY09595	APAR IY10564
7017-S80	No	Yes	Yes
7017-S80	Yes	Yes	Yes
4/16 token ring (# 4959)	Yes	Yes	Yes
Dual Ultra2 SCSI PCI (# 6205)	Yes	Yes	Yes
2104-DL1/TL1	Yes	Yes	Yes
Gigabit Fibre Channel ¹ (# 6227)	No	Yes	Yes
2102-F10 ¹ 2102-D00 ¹	No	Yes	Yes
Advanced SerialRAID Plus (# 6230)	Yes	Yes	Yes
32 MB Fast Write Cache (# 6235) (2-initiator only)	Yes	Yes	Yes
7044-170	No	Yes	Yes
7044-270	No	Yes	Yes
Ultra SCSI F/W # 6204	No	Yes	Yes
4-Port 10/100 Ethernet PCI Adapter (# 4951)	No	Yes	Yes
¹ Not supported in HA SP environments.			

Chapter 17. Parallel Software

This chapter covers the major parallel software products installed principally for systems management and application development in an RS/6000 SP clustered environment.

The products that will be covered are:

- Parallel System Support Programs for AIX (PSSP) - (5765-D51)
- General Parallel File System (GPFS) - (5765-B95)
- LoadLeveler Version 2.1.0 - (5765-D61)
- Parallel Environment Version 2.4 - (5765-543)
- Parallel Engineering and Scientific Subroutine Library (ESSL) - (5765-C42)
- Parallel Optimization Subroutine Library (OSL) - (5765-392)

17.1 Parallel System Support Programs for AIX (PSSP)

PSSP (5765-D51) is a collection of administrative and operational software applications that run on each node of an RS/6000 SP as well as on SP-attached servers. Built upon the system management tools and commands of the AIX, Version 4, operating system, PSSP enables system administrators and operators to better manage SP systems and their environments.

Sets of software tools and related utilities, including application programming interfaces (APIs), have been grouped together to offer easier administration of installation, configuration, device management, security administration, error logging, system recovery, and resource accounting in the SP environment.

The following are the highlights of PSSP:

- A full suite of system management applications with the unique functions required to manage the RS/6000 SP and the SP-attached servers.
- Simplified installation, operation, and maintenance of all nodes in an SP system, from a single Control Workstation.
- Advanced error detection and recovery features that reduce the impact and occurrence of unplanned outages.
- Parallel system management tools for allocating SP resources across the enterprise.

- Coexistence of up to three releases within an SP partition allowing for easier software migration.
- Advanced performance monitoring for consolidated analysis and reporting.

17.1.1 PSSP Features and Benefits

The following items are the major features and benefits of PSSP:

- Single point of control
Allows system administrator or operator to perform all local and remote administrative functions from the Control Workstation.
- RS/6000 SP Perspectives
Provides easier access to SP subsystems through an object-oriented management interface (GUI).
Allows monitoring and controlling hardware, creating and monitoring events, defining and managing virtual shared disks, and configuring systems.
- Consolidated accounting
Enhances AIX node data accounting and consolidates information with summaries by node classes.
Provides accounting hooks to charge for exclusive node use.
- System partitioning
Allows testing of different levels and PTFs of operating system, system software, or hardware on subsets of nodes.
Allows different production environments to be concurrently executed for work load isolation.
- System monitoring and control
Enables the system administrator to gracefully shut down, rather than reboot, nodes or the complete system.
Allows authorized users to monitor and manipulate SP hardware variables at node or frame levels.
Provides for consolidation of error and status logs for expedited problem determination.
- RS/6000 cluster technology
Provides improved infrastructures for event monitoring and recovery coordination.

API for building highly available distributed applications.

- Node isolation
Allows non-disruptive removal or retrieval of nodes from SP configuration without a switch fault.
- HACWS
Allows all SP administrative functions to continue uninterrupted in the event of a failure of the control workstation.
Allows for automatic backup of primary control workstation using HACMP.
Provides for backup access to SDR.
- Performance Toolbox Parallel Extensions (PTPE) (5648-B64)
Collects and displays statistical data on SP hardware and software and monitors system performance.
Simplifies run-time performance monitoring on a large number of nodes.
- Virtual Shared Disk (VSD)
Allows multiple nodes to access a read disk as if the disk were attached locally to each node.
- Recoverable VSD (5765-646)
Allows transparent recovery of VSD.
- Migration and coexistence
Allows up to two additional levels of PSSP (2.2, 2.3, 2.4) to run with PSSP 3.1 in a single SP partition.

17.1.2 PSSP Prerequisites

The following are the hardware and software prerequisites for PSSP:

- Hardware supported: RS/6000 SP, S80, S70 Advanced
- Operating systems: AIX Version 4.3.2 or later
- User interfaces: SP Perspectives GUI, command line
- Prerequisite software: C for AIX Version 4.3
- Prerequisite for PTPE: Performance Toolbox for AIX Version 2.2 (5765-654)
- Prerequisite for HACWS: HACMP Version 4.3 for AIX (5765-D28)

17.1.3 RS/6000 Cluster Technology

The PSSP's RS/6000 Cluster Technology is a collection of services that define hardware and software resources, node relationships and coordinated actions to manage groups of nodes on an SP.

Event Management monitors hardware and software resources in the SP and notifies an interested application or subsystem when the state of the resource changes. Resources on any node in the SP can be monitored from any other node.

Topology Services define the relationships between nodes in a cluster to allow seamless takeover of functions in the event of a node failure.

Group Services provides a set of interfaces that enable a distributed subsystem, such as General Parallel File System (GPFS), to synchronize recovery actions among the processes making up the subsystem.

17.1.4 Subsystem Communications Support

The Communication Subsystems Support component contains SP Switch adapter diagnostics, switch initialization and fault-handling software, device driver and configuration methods (configure and unconfigure), as well as parallel communications APIs.

17.1.5 Enabling Parallel Processing

The Virtual Shared Disk (VSD) API component creates logical disk volumes for parallel application access of a real disk device. These can be attached locally or to another SP node. This feature can enhance the performance of applications that provide concurrency control for data integrity, such as Oracle databases.

The IBM Recoverable Virtual Shared Disk function provides recovery from failures of virtual shared disk server nodes and takes advantage of the availability services provided by PSSP to determine which nodes are up and operational.

17.1.6 Application Programming Models

PSSP supports a multithreaded, standards-compliant Message Passing Interface (MPI) using IBM Parallel Environment for AIX (PE), as well as maintaining its single-threaded MPI support. In addition, PSSP includes a Low-level Application Programming Interface (LAPI) with a flexible, active message style, and communications programming model on the SP Switch.

17.1.7 Collecting Performance Data

The Performance Toolbox Parallel Extensions (PTPE) function of PSSP collects and provides performance data for SP hardware and software through enhancements to the Performance Toolbox for AIX (PTX) product, the preferred performance monitor for AIX systems. It allows PTX to monitor unique SP subsystems, such as VSD, SP Switch, and LoadLeveler.

PTPE organizes the SP into a set of performance reporting groups with coordinating managers and distributes the burden of monitoring nodes throughout the SP system, thus eliminating the need for a dedicated monitoring node. PTPE also provides average performance statistics for the SP system rather than monitoring every data point on all SP nodes. This can help reduce the computational effort required for run-time monitoring of SP performance.

17.1.8 Other Available Services Included

PSSP also includes the following publicly available software packages:

- Network Time Protocol (NTP) for clock synchronization across SP nodes
- Perl programming language for developing system-wide shell scripts
- Software Update Protocol (SUP) for installing software from the boot file server
- Kerberos IV security for authentication of the execution of remote commands
- Tool command language (Tcl) for controlling and extending applications

17.2 General Parallel File System (GPFS)

GPFS (5765-B95) is a standards-based, parallel file system that delivers high performance, high availability, and high scalability while preserving the application interfaces used in standard file systems. GPFS allows access to files within an SP system from any GPFS node and can be exploited by parallel jobs running on multiple SP nodes as well as by serial applications that are scheduled to nodes based on processor availability.

IBM originally introduced Parallel Input Output File System (PIOFS) to speed file access on the SP, but PIOFS is limited in usage due to its lack of recoverability. GPFS is designed expressly to deliver scalable performance across multiple file system nodes, to comply with UNIX file standards, and to be recoverable from most failures. It delivers all of the functions covered in the following sections.

17.2.1 Administration

GPFS provides functions that simplify multi-node administration and can be performed from any node in the SP configuration. These functions are based on, and are in addition to, the AIX administrative commands that continue to operate. A single GPFS multi-node command can perform a file system function across the entire SP system. In addition, most existing UNIX utilities will also run unchanged. All of these capabilities allow GPFS to be used as a replacement for existing UNIX file systems where parallel optimization is desired.

17.2.2 Standards Compliance

GPFS supports the file system standards of X/Open Version 4.0., with minor exceptions, allowing most AIX and UNIX applications to use GPFS data without modification.

17.2.3 Higher Performance and Scalability

By delivering file performance across multiple nodes and disks, GPFS scales beyond single-server (node) performance limits. This level of performance is achieved through the use of IBM Virtual Shared Disk, client-side data caching, large file block support, and the ability to perform read-ahead and write-behind file functions. As a result, GPFS can outperform PIOFS, Network File System (NFS), Distributed File System (DFS), and Journaled File System (JFS) in a parallel environment. Unlike NFS and JFS, GPFS file performance scales as additional file server nodes and disks are added to the SP system.

17.2.4 Availability and Recoverability

GPFS can survive many system and I/O failures. Through its use of the RS/6000 Cluster Technology capabilities of IBM Parallel System Support Programs for AIX, in combination with IBM Recoverable Virtual Shared Disk (RVSD), GPFS is able to automatically recover from node, disk connection, and disk adapter failures. GPFS will transparently failover lock servers and other GPFS central services. Through its use of RVSD, GPFS continues to operate in the event of disk connection failures. GPFS allows data replication to further reduce the chances of losing data if the storage media fails. Unlike PIOFS, GPFS is a logging file system that allows the re-creation of consistent structures for quicker recovery after node failures. GPFS also provides the capability to mount multiple file systems, each of which can have its own recovery scope in the event of component failures.

17.3 LoadLeveler

LoadLeveler, Version 2.1.0 (5765-D61), is a distributed network-wide job-management program for dynamically scheduling work on IBM and non-IBM processors. LoadLeveler balances your work load by efficiently managing job flow in the network and distributing work across all LoadLeveler-defined hardware resources.

In a network of UNIX workstations or RS/6000 SP systems, LoadLeveler matches processing needs to available resources for improved performance and faster turnaround. LoadLeveler is an application that runs as a set of background processes (daemons) on each IBM RS/6000 workstation and SP node in the network. It provides a facility for building, submitting, and processing batch jobs quickly and efficiently in a dynamic environment with a non-predictive job arrival pattern. For greater efficiency on an SP system, LoadLeveler with the Communication Subsystems Support (CSS) function of PSSP Version 3.1 can support up to four user space tasks per SP Switch adapter.

17.3.1 LoadLeveler Highlights

The major functions of LoadLeveler can be expressed as follows:

- Distributed, full-function job scheduler
- Serial and parallel batch and interactive workload balancing
- Central point-of-control for workload administration
- Full scalability across processors and jobs
- API to enable alternate scheduling algorithms
- High availability that provides automatic recovery of central scheduler and can be configured with HACMP for node and network failover

17.3.2 LoadLeveler Interfaces

LoadLeveler has a command line interface and a Motif-based graphical user interface (GUI), making it easy to submit and cancel jobs, monitor job status, and set and change job priorities. The system can be configured at the network and node levels, where workstations or SP nodes may be identified as job submitters, compute servers, or both. When a job is scheduled, its requirements are compared with all the resources known to LoadLeveler. Job requirements might be a combination of memory, disk space, architecture, operating system, and application programs. LoadLeveler's Central Manager collects resource information and dispatches the job as soon as it locates suitable nodes.

LoadLeveler offers the option of using its own scheduler or an API to use alternative schedulers, such as EASY from the Cornell Theory Center. Details on EASY can be obtained from the Cornell Theory Center Web site:

<http://www.tc.cornell.edu/>

LoadLeveler also provides a user-initiated or system-initiated checkpoint/restart capability for certain types of Fortran, C, or C++ jobs linked to the LoadLeveler libraries.

17.3.3 What to Expect from LoadLeveler

For parallel jobs, LoadLeveler interfaces with parallel programming software, Parallel Environment for AIX or PVM, to obtain the multiple SP nodes required for the job's parallel tasks. APIs are also available for linking to other parallel application environments. LoadLeveler Version 2.1 allows a task to run both the Message Passing Interface (MPI) and the Low-level Application Programming Interface (LAPI) communications protocols.

- Individual Control

At the node level, users can specify to LoadLeveler when their processing nodes are available and how they are to be used. For example, some users might let their workstations accept any job during the night, but only certain jobs during the day, when they most need their resources. Other users might simply tell LoadLeveler to monitor their keyboard activity and make their workstations available whenever they've been idle for a sufficient time.

- Central Control

From a system management perspective, LoadLeveler allows a system administrator to control all jobs running in a cluster including the SP system. With hundreds of nodes configured, job and machine status are always available, providing administrators with the information needed to make adjustments to job classes and changes to LoadLeveler-controlled resources.

- Scalability

As nodes are added, LoadLeveler automatically scales upward so that the additional resources are transparent to the user.

17.4 Parallel Environment

Parallel Environment Version 2.4 (5765-543) is a complete solution for enterprises that need to develop, debug, analyze, tune, and execute parallel programs on the AIX platform.

17.4.1 Key Elements of PE

This application development solution consists of:

- Parallel Message Passing APIs for full implementation of the MPI 1.2 standard plus a subset of the MPI 1.2 I/O chapter of the MPI 1.2 standard.
- Continued support for IBM's own Message Passing Library for communications between executing tasks in a Fortran, C, or C++ parallel program.
- A Parallel Operating Environment (POE) for managing the development and execution of parallel applications.
- A user-friendly visualization and performance monitoring tool, Visualization Tool (VT), for viewing the unique performance characteristics of a parallel application. The VT graphically shows the performance and communication characteristics of an application and can act as an online monitor to show current application activity.
- Parallel debuggers offering both command-line and Motif-based interfaces. These debuggers extend traditional AIX capabilities and provide features for parallel application task debugging.

17.4.2 PE Highlights

The following are the major highlights of Parallel Environment:

- Provides a full-function development environment for parallel applications
- Supports AIX Version 4.3.2 and PSSP Version 3.1
- Exploits threads and thread-safe MPI message passing on all nodes including symmetric multiprocessors (SMPs)
- Supports up to four user processes per SP node adapter
- Supports Low-level Application Programming Interface (LAPI) programs
- Easier parallel application development
- Enhanced support for threaded and LAPI applications
- Improved profiling for analysis of application performance
- Enhanced XProfiler graphical performance tool
- Enhanced visualization and monitoring tools for threaded applications
- Enables easy application portability to networked RS/6000 or RS/6000 SP systems

17.4.3 Parallel Environment Functions

Parallel Environment (PE) provides the following functions:

- PE provides exploitation of threads and thread-safe MPI message passing on all nodes, including SMPs, and also supports Low-level Application Programming Interface (LAPI) programs. Included are: Tool extension for threaded applications, support for parallel applications on AIX 4.2.1, and easy application portability to a networked cluster of RS/6000 or RS/6000 SP systems.
- The POE supports threaded applications and the handling of asynchronous signals from the user parallel program. The visualization and performance monitoring tool has been enhanced to support thread-safe VT trace output files in addition to limited visualizations of thread events. A graphical tool, XProfiler, is available for analyzing application performance. Zoom, filter, and search functions support graphical manipulation to aid analysis and tuning of parallel programs.
- The number of tasks supported in a single POE job is 2048 (1024 with user space MPI or LAPI libraries). In addition, up to four user space tasks can now run on each SP node from the same or different jobs, therefore, allowing parallel applications to exploit SMP nodes.
- POE also provides the ability to checkpoint the state of a parallel batch program for restart in case of application failure.
- For increased MPI application debugging, the PEDB capability provides the following:
 - A summary of the number of active messages from each task in an application
 - Message queue information for a specific task
 - Detailed confirmation for a specific message
 - The ability to debug parallel applications currently executing
 - The ability to view thread events in the visualization and performance monitoring tool

17.4.4 Software Requirements

PE Version 2.4 requires AIX Version 4.3.2 and PSSP Version 3.1.

17.5 Parallel Engineering and Scientific Subroutine Library (ESSL)

Parallel ESSL (5765-C42) Version 3, is a collection of mathematical subroutines that provides a wide range of high-performance, mathematical

functions for many different scientific and engineering applications. Parallel ESSL is specifically tuned to exploit the full power of the SP hardware with scalability across the range of system configurations. In addition to SP systems, Parallel ESSL runs on clusters of RS/6000 servers and workstations.

Parallel ESSL can be used to develop and enable many different types of scientific and engineering applications. New applications can be designed to take advantage of all the capabilities of ESSL. Existing applications can be easily enabled by replacing comparable routines and inline code with calls to ESSL subroutines.

17.5.1 Parallel ESSL Highlights

The following lists some highlights of Parallel ESSL:

- Provides mathematical algorithms optimized for high performance
- Supports SP systems and clusters of RS/6000 servers and workstations
- Callable from XL Fortran, XL High Performance Fortran (HPF), C, and C Set++ applications
- High performance - Designed for high mathematical computational performance
- Versatile - Supports many scientific and engineering applications used by multiple industries
- RS/6000 compatible - Tuned to the characteristics of RS/6000 hardware
- Ease of use - Can be used with existing programs through relinking rather than recompiling
- Easy application development - Supports easy development of parallel applications with the Single Program Multiple Data model and the Shared Memory Parallel Processing model for SMPs

17.5.2 Wide Range of Mathematical Functions

ESSL provides a variety of complex mathematical functions, such as:

- Basic Linear Algebra Subroutines (BLAS)
- Linear Algebraic Equations
- Eigensystem Analysis
- Fourier Transforms

17.5.3 Application Environments

Examples of applications that use these types of mathematical subroutines are:

- Structural analysis
- Time series analysis
- Computational chemistry
- Computational techniques
- Fluid dynamics analysis
- Mathematical analysis
- Seismic analysis
- Dynamic systems simulation
- Reservoir modeling
- Nuclear engineering
- Quantitative analysis
- Electronic circuit design

17.5.4 Compatibility with Other Subroutine Libraries

The ESSL products are compatible with public domain subroutine libraries such as Basic Linear Algebra Subprograms (BLAS), Scalable Linear Algebra Package (ScaLAPACK), and Parallel Basic Linear Algebra Subprograms (PBLAS), thus making it easy to migrate applications that utilize these libraries to Parallel ESSL.

Also new with Parallel ESSL Version 2.1 is support for additional subroutines: Fortran90 and Fortran77 iterative solvers for general sparse matrices.

17.5.5 SMP Support

Parallel ESSL, with its thread-safe subroutines, supports SMPs. Key subroutines are multithreaded for enhanced SMP performance. This version of ESSL also provides the ability to dynamically allocate storage for subroutines that require extra working storage to perform computations.

Version 2.1 of Parallel ESSL for AIX runs on stand-alone RS/6000 SMP models as well as on SMP nodes in SP systems, therefore, providing both SMP and thread-tolerant POWER2 libraries. Support is provided for both the Message Passing Interface (MPI) threaded library and the MPI signal-handling library.

Parallel ESSL is specifically tuned to RS/6000 hardware for increased performance. To take advantage of the increased performance on SMP systems, programs need only be relinked, not recompiled.

17.5.6 Useful URLs

The following Web sites provide additional information on Parallel ESSL:

www.rs6000.ibm.com/resource/aix_resource/sp_books/essl/index.html

www.rs6000.ibm.com/software/Apps/essl.html

www.rs6000.ibm.com/software/sp_products/esslpara.html

17.6 Parallel Optimization Subroutine Library (OSL)

Parallel OSL (5765-392) mathematical optimization on the SP is a collection of high-performance, mathematical subroutines used by application programmers to solve large optimization problems.

Mathematical programming techniques can be applied to problems where the user wants to minimize or maximize an objective function, subject to a set of constraints. A feasible solution solves the constraints; an optimal solution yields the largest or smallest value of the objective function among all feasible solutions.

Parallel OSL is a set of over 60 subroutines callable from application programs to find the optimal solution to several types of problems using linear programming (LP), mixed-integer programming (MIP), and quadratic programming (QP) mathematical techniques. Some of the solutions use serial algorithms; that is, all computations are performed in sequence on an RS/6000 or on a single node of an SP processor. Others use algorithms that exploit the parallel processing capabilities of the SP system so that multiple nodes may concurrently perform computations on subtasks of the problem to be solved.

Parallel OSL models may be defined using OSL data structures. Data structures include the Mathematical Programming System (MPS) format for compatibility with predecessor IBM mathematical programming products or the Lotus 1-2-3 spreadsheet format. Subroutines are provided for loading models in any of these formats.

17.6.1 Parallel Application Programming

Only minor changes to serial OSL applications are required to generate parallel application programs. No explicit parallel programming is required. The parallel solvers have the same names and calling sequences as their serial counterparts.

Parallel OSL's mathematical subroutines are callable from user application programs written in XL Fortran or C. High-level subroutines can solve a problem with the user having minimal knowledge of mathematical programming. Low-level subroutines give the user the flexibility to structure algorithms without having to write new routines independently.

A supported parallel execution environment, such as Parallel Environment, Version 2.4 (5765-543), must be present to execute Parallel OSL on concurrent parallel processors. This provides the environment for the execution in parallel of subproblems of the mathematical optimization, including allocation of multiple processors and invocation and monitoring of subproblem execution.

17.6.2 Linear Programming (LP)

In an LP algorithm, both the objective function and the constraints are linear. The algorithms available in Parallel OSL are:

- Simplex method - OSLp uses either a primal or a dual simplex serial algorithm.
- Interior Point Barrier method - OSLp uses the primal barrier, primal-dual barrier or primal-dual with predictor-corrector, fine-grained parallelized algorithms.
- Network Solver method - OSLp solves this special case using a serial algorithm.

17.6.3 Mixed-Integer Programming (MIP)

MIP problems are LPs in which some variables are constrained to be integers. Parallel OSL includes a versatile, course-grain, parallelized branch-and-bound solver, handling MIP problems with either linear or quadratic objective functions. The simplex solver (primal or dual, selectable) is used on LP sub-problems; the QP solver is used on QP sub-problems. As many SP nodes as are available (and licensed for OSLp) may be used to process different branches of the search tree and improve the performance of the algorithm.

17.6.4 Quadratic Programming (QP)

QP problems have a convex quadratic objective function and linear constraints. Parallel OSL includes a fast two-stage serial algorithm for solving these problems. The first sub-algorithm solves an approximating LP problem and a related, very simple QP problem at each iteration. When successive approximations are close enough together, the second sub-algorithm is used. This extension of the simplex method permits a quadratic objective function and converges very rapidly when given a good starting value.

17.6.5 Other Solution Capabilities

Subroutines are provided for sensitivity analysis, allowing evaluation of the effect of objective function coefficient changes on the optimal solution basis. Also provided is a parametric solver that shows how the objective value and optimal solution vary as row bounds, column bounds, and objective function coefficients change over a range.

17.6.6 Parallel OSL Performance

Benchmarks have shown that Parallel OSL achieves significant speed-ups on LP and MIP problems with its parallelized solution algorithms. LPs solved with the interior point algorithm are typically able to achieve a sub-linear speed-up of 40 to 50 percent processor utilization. That is, running such a problem on an SP with eight processors is three times as fast as running OSL/6000 on a single node.

MIP problems perform even better on a parallel processor because the MIP branch-and-bound solution algorithm requires little interprocess communication; so, multiple processors can operate concurrently and independently with little need for synchronization. Nearly linear speed-up is regularly achieved with occasional super-linear speed-up. Processor utilization is typically in the 90 to 100 percent range, resulting in a speed-up of seven to eight times on an eight-processor SP.

17.6.7 Parallel OSL Prerequisites

The following are the required levels of hardware and software:

- Hardware supported: SP, SP1, RS/6000, RS/6000 cluster
- Operating system: AIX Version 4.3
- Parallel environment: Parallel Environment for AIX Version 2.4 (5765-543)
- User interfaces: Callable from user application drivers

- Languages supported: AIX XL Fortran Compiler/6000 Version 2.3 and Version 3.1, XL C Compiler Version 1.3, AIX XL C++ Compiler/6000 Version 1.1.

Chapter 18. CATIA Solutions

The CATIA Solutions product line offers customers a wide range of easy-to-use, scalable computer-aided design (CAD), computer-aided manufacturing (CAM), and computer-aided engineering (CAE) solutions. These applications cover the entire product lifecycle from initial concept through detailed design, refinement, testing, and manufacture. The RS/6000 Operating Environment for CATIA solution facilitates the installation of CATIA and ENOVIA CATweb solutions on RS/6000 AIX systems.

18.1 CATIA Solutions Overview

CATIA Solutions provide:

- Process-oriented solutions to digitally design, analyze, optimize and manufacture products in a wide range of industries
- An enterprise-wide portal to digital product data and advanced tools to control the entire product life-cycle

The CATIA and ENOVIA product families from IBM Product Lifecycle Management (PLM) include a full range of technology and services that can empower industries worldwide with a seamless and complete leading-edge product development environment.

- Aerospace

From civilian and military aircraft to satellites and launch vehicles, PLM serves all segments of the world's aerospace and defense companies. In an industry where reliability isn't optional, the majority of the world's aviation companies rely on CATIA and other engineering solutions from IBM.

- Automotive

As a major provider to automotive companies of all sizes around the world, PLM offers market-proven technology and services for the development of passenger cars and light trucks, commercial trucks and buses, racing cars, motorcycles and specialty vehicles.

- Consumer Goods

Digital mock-ups, virtual prototypes and intra-company product data management are just a few of the ways that consumer goods manufacturers gain advantage from the wide range of products offered by PLM.

- Electrical and Electronics

Original equipment manufacturers of electrical and electronic equipment and their suppliers benefit from IBM PLM's integrated technology and services. The ability to tightly integrate the tremendous advancements in ECAD with CATIA and ENOVIA's advanced MCAD and PDM solutions can give electrical and electronics companies a significant competitive edge.

- Industrial Machinery

IBM PLM supports companies across the industrial machinery industry, including original equipment manufacturers of heavy, construction, textile, agricultural and paper machinery.

- AEC Plant Design

The AEC plant design industry benefits from IBM PLM's industry-leading CAD/CAM/CAE and product data management software and services.

- Shipbuilding

IBM PLM serves manufacturers and operators of cruise ships, commercial and specialized vessels, surface combatants and submarines by facilitating and automating critical product development steps: concept; hull surfacing and fairing, outfitting, production and commissioning.

- Additional Industries

IBM PLM supports companies in almost every conceivable industry to become more competitive by reducing development costs and time-to-market, while improving innovation and quality.

For more information on CATIA and CADAM, visit the following Web site:

www.ibm.com/solutions/plm

18.2 Operating Environment for CATIA

The RS/6000 Operating Environment for CATIA, V2.3 helps automate the installation of a CATIA or ENOVIA CATweb solution on AIX Version 4.3.3.

Specific levels of CATIA supported with this tool are:

- CATIA 5.3
- CATIA 4.2.2 Refresh 1
- CATIA 4.2.1 Refresh 2
- CATIA 4.2.1 Refresh 1
- CATIA 4.2.0 Refresh 1

- ENOVIA CATweb Solutions, Version 2.2

Dassault Systemes products and the AIX operating system are not included in this solution and must be purchased separately.

The IBM RS/6000 Operating Environment for CATIA, Version 2.3, consists of these programs:

- RS/6000 CATIA/CATweb Installation Utility
- RS/6000 CATIA Welcome Center
- IBM XL FORTRAN Runtime Environment for AIX, Version 5.1 (5765-C11)
- PTFs at CASIL 0003 level (including support for the RS/6000 44P)
- Device Drivers (including Spaceball* and Magellan*) and X11R6
- IBM HTTP Server Version 1.3.6, international version
- Adobe Acrobat, Version 3.0
- Netscape Communicator 4.7, international version

In addition to the base products, the RS/6000 CATIA/CATweb Installation Utility provides installation and configuration assistance for the following products (product license and media required):

- IBM XL FORTRAN for AIX, Version 5.1 (5765-C10)
- IBM C and C++ Compilers for OS/2, AIX, and Windows NT, Version 6.6
- IBM DB2 Universal Database Enterprise Edition Version 6.1 (5648-B97)
- IBM DB2 Data Links Manager, Version 6.1 (5648-B90)

The IBM IDSICC lab has tested CATIA and CATweb with the levels of software integrated in this solution. The installation utility, Quick Start Installation Guide, and 3D input device drivers help simplify the installation of a CATIA or CATweb operating environment. The RS/6000 CATIA Welcome Center provides a user-friendly environment for Internet access, product information, and registration. Netscape Communicator, IBM HTTP Server, and Adobe Acrobat Reader enable Web communication.

IBM HTTP Server and Netscape Communicator (international version 56-bit DES encryption) are included with this solution as default Web enablement for a ENOVIA CATweb Solutions installation. The RS/6000 CATIA/CATweb Installation Utility allows you to override the default and install a more secure version of these products on the AIX, Version 4.3 Bonus Pack (a more secure North American version is 128-bit DES encryption).

The powerful AIX OpenGL and X11 graphics libraries enable ENOVIA CATweb Solutions customers to take advantage of multiple processors for near linear scalability or utilize Virtual Frame Buffer (VFB) to run ENOVIA CATweb Solutions on an RS/6000 server without the need for a graphics adapter.

Appendix A. HACMP History

This appendix provides a historical view of HACMP.

A.1 HACMP Version 1.1

This first version of HACMP, released June, 1992, stated that it aimed to bring together cluster processing and high availability to open system client/server configurations. It included utilities for the *no single point of failure* concept that underpins HACMP.

The following statement is taken from the IBM announcement letter, ZP92-0325, as is, because it succinctly described the purpose of the software.

HACMP/6000 software controls the cluster environment behavior. It is also a reporting and management tool for system administrators or programmers to use in implementing applications for HACMP/6000 cluster systems or in managing HACMP/6000 systems configurations where mission critical application failover is a requirement.

In a simple failover configuration, the cluster manager provides for prompt restart of an application or subsystem on a backup RISC System/6000 processor after there has been a failure of the primary processor. The backup processor can be a dedicated stand-by or actively operational in a load sharing configuration with the primary processor.

There were three cluster configurations known as mode 1, 2, and 3, although only mode 1 and 2 were available with version 1.1.

- Mode 1 - Standby mode, where each cluster has an active processor and a standby processor.
- Mode 2 - Partitioned workload, where each cluster supports different users and can permit each of the two processors to be protected by the other.

The HAS and CRM features comprised version 1.1.

A.1.1 HACMP Version 1.2

HACMP Version 1.2 was released in May, 1993 and included a management information base (MIB) for the simple network management protocol (SNMP) and allowed for a cluster to be monitored from a network management console. This introduced the SNMP sub-agent, clsmuxpd.

- Monitoring of the cluster was further improved as AIX tracing was now supported for clstrmgr, clinfo, cclockd, and clsmuxpd daemons.
- HACMP Version 1.2 was extended to use the standard AIX error notification facility.
- Mode 3 configurations became available in version 1.2.

Mode 3 - Loosely coupled multiprocessing, which allowed two processors to work on the same shared workload.

A.2 HACMP Version 2.1

In this version of HACMP, released in December, 1993, system management was done through SMIT. It was possible to install and configure a cluster from a single processor. Cluster diagnostics and verification of hardware, software, and configuration levels for all nodes in the cluster were included in the product.

- Resources could be disk, volume groups, file systems, networks, or applications.
- Cluster configurations were now described in the terms used today.
 - Standby configurations. These were the traditional redundant hardware configurations where one or more standby nodes stands idle.
 - Takeover configurations. In these configurations, all cluster nodes were doing useful work.

Standby and takeover configurations could handle resources in one of the following ways:

- Rotating failover to a standby machine(s)
- Standby or takeover configurations with cascading resource groups
- Concurrent access
- Cluster customizing was enhanced by supporting user pre- and post-event scripts.
- Concurrent LVM and concurrent access subsystems extended for new disk subsystems (IBM 9333 and IBM 7135-110 Radiant array).

A.3 HACMP Version 3.1

HACMP Version 3.1 was released in December, 1994. Even at this early stage in the product's development, scalability was an issue, and support for eight RS/6000 nodes in a cluster was added.

- Resource groups were introduced to logically group resources, thus improving an administrators ability to manage resources.
- Similarly, cascading failover, allowing more than one node to be configured to take over a given resource group, was an additional configuration option introduced in Version 3.1.
- There were changes made to the algorithm used for the keep-alive subsystem reducing network traffic and an option in SMIT to allow the administrator to modify the time used to detect when a heartbeat path had failed and, therefore, initiate failover sooner.

A.3.1 HACMP Version 3.1.1

With version 3.1.1, released in December, 1994, HACMP became available on the RS/6000 SP, which increased the application base that might make use of high availability. The SP has a very fast communication subsystem that allows large quantities of data to be transferred quickly between nodes. Customers now had high availability on two different architectures, stand-alone RS/6000s and the RS/6000 SP.

Support for eight nodes in a cluster on an SP and node failover using the SP high performance switch (HPS) was added.

A.4 HACMP Version 4.1

Important for cluster administrators, HACMP Version 4.1, released July, 1995, includes version compatibility to allow a cluster to be upgraded from an earlier software release without taking the whole cluster off-line.

- The CRM feature optionally adds concurrent shared access management for supported disks.
- Support is added for IBMs symmetric multiprocessor (SMP) machines.

A.4.1 HACMP Version 4.1.1

The following sections introduce specific HACMP levels:

- New features were announced with Version 4.1.1 of HACMP in May, 1996, namely HANFS, which enhances the availability of data accessed by NFS, and as a separate product HAGEO, which supports HACMP over WAN networks.
- HACMP Version 4.1.1 also included the introduction of the cluster snapshot utility, which is used extensively in the management of clusters to clone or recover a cluster and also for problem determination. Using the

cluster snapshot utility, the current cluster configuration can be captured as two simple ASCII files.

- There was a further enhancement to concurrent access to support serial link and serial storage architecture (SSA).

A.4.2 HACMP Version 4.2.0

Dynamic Automatic Reconfiguration Events (DARE) was introduced in July, 1996 and allows clusters to be modified while up and running. The cluster topology can be changed, a node removed or added to the cluster, network adapters and networks modified, even failover behavior changed while the cluster is active. Single dynamic reconfiguration events can be used to modify resource groups.

The cluster single point of control (C-SPOC) allows resources to be modified by any node in the cluster regardless of which node currently controls the resource. At this time, C-SPOC is only available for a two node cluster.

C-SPOC commands can be used to administer user accounts and make changes to shared LVM components taking advantage of advances in the underlying logical volume manager (LVM).

Changes resulting from C-SPOC commands do not have to be manually resynchronized across the cluster. This feature is known as *lazy update* and was a significant step forward in reducing downtime for a cluster.

This small section does not do justice to the benefit system administrators derive from these features.

A.4.3 HACMP Version 4.2.1

In HACMP, Version 4.2.1, introduced in May, 1997, HAView is now packaged with HACMP and extends the TME 10 NetView product to provide monitoring of clusters and cluster components across a network using a graphical interface.

- Enhanced security, which uses kerberos authentication, is now supported with HACMP on the RS/6000 SP.
- HAGEO is supported in version 4.2.1 while HACMP, Version 4.2.0, does not support HAGEO.

A.4.4 HACMP ES Version 4.2.1

The first release of HACMP ES, Version 4.2.1, was announced shortly after HACMP 4.2.1 in September, 1997. It can only be installed on SP nodes and was scalable beyond the eight node limit that exists in HACMP.

- Up to 16 RS/6000 SP nodes in a single cluster are supported with HACMP ES, Version 4.2.1.
- The heart beating used by the HACMP ES cluster manager exploits the cluster technology in the parallel system support program (PSSP) on the SP, specifically Event Management, group services, and topology services.

It was then possible to define new events using the facilities provided, at that time, in PSSP 2.3 and to define customer recovery programs to be run upon detection of such events.

- The ES version of HACMP now uses the same configuration features as the HACMP options; so, migration is possible from HACMP to HACMP ES without losing any custom configurations.

There are, however, some functional limitations of HACMP ES, Version 4.2.1, which are listed as follows:

- Fast failover is not supported.
- C-SPOC is not supported with a cluster greater than eight nodes.
- VSM is not supported with a cluster greater than eight nodes.
- Enhanced Scalability cannot execute in the same cluster as the HAS, CRM, or HANFS for AIX features.
- Concurrent access configurations are not supported.
- Cluster Lock Manager is not supported.
- Dynamic reconfiguration of topology is not supported.
- It is necessary to upgrade all nodes on the cluster at the same time.
- ATM, SOCC, and slip networks are not supported.

As this overview continues, you will see that the functional limitations of HACMP ES have been addressed. Customers who opted for HACMP ES identified the benefits as being derived from being able to use the facilities in PSSP.

A.4.5 HACMP and HACMP ES Version 4.2.2

Fast recovery for systems with large and complex disk configurations was introduced in October, 1997. Faster failover is achieved by running recovery phases simultaneously as opposed to serially as they had been in the past. Also, the option was added to run full fsck or logredo for a file system.

- A system administrator using version 4.2.2 can evaluate new event scripts with event emulation. Events are now emulated without affecting the cluster and the result reported.
- Dynamic resource movement is possible with the introduction of the `clhare` utility.
- There are significant enhancements to the verification process and the `clverify` command to check name resolution for IP labels used in the cluster configuration, the network conflicts where hardware address takeover is configured, the `lvm`, and the rotating resource groups.

In addition, `clverify` can now support custom verification methods.

- Support for NFS Version 3 is now included.
- There are still limitations in HACMP ES as there were for version 4.2.1.
- In March, 1998, target mode SSA was announced for HACMP, Version 4.2.2 only.

A.4.6 HACMP and HACMP ES Version 4.3

Custom snapshot methods are introduced in HACMP Version 4.3.0, in September, 1998, allowing an administrator to capture additional information with a snapshot. The ability to capture a cluster snapshot was added to the HACMP ES feature.

- C-SPOC support for shared volumes and users or groups is further extended. All nodes in a cluster are now immediately informed of any changes; so, failover time is reduced, as *lazy update* will usually not need to be used.
- A task guide is introduced for shared volume groups that aims to reduce the time spent on some administration tasks within the cluster (`smit hacmp -> Cluster System Management -> Task Guide For Creating a Shared Volume Group`).
- Multiple pre- and post-event scripts are now supported in version 4.3.0, extending once more the level to which HACMP or HACMP ES cluster events can be tailored.

- HACMP ES is now available on stand-alone RS/6000s. RISC system cluster technology (RSCT), previously packaged with PSSP, is now packaged with HACMP ES.
- Scalability in HACMP ES is increased once more to support up to 32 nodes in an SP cluster. Stand-alone systems are still limited to eight nodes.
- The CRM feature was added to HACMP ES. Eight nodes in a cluster are supported for ESCRM.

Note

In October, 1998, VSS support was announced for HACMP Versions 4.1.1, 4.2.2, and 4.3.

A.4.7 HACMP and HACMP ES Version 4.3.1

Support for target mode SSA is added to the HACMP ES and ESCRM features with version 4.3.1 in July, 1999.

- HACMP, Version 4.3.1, supports the AIX Fast Connect application as a highly-available resource, and the clverify utility is enhanced to detect errors in the configuration of AIX Fast Connect with HACMP.
- The dynamic reconfiguration resource migration utility is now accessible through SMIT panels as well as the original command line interface, therefore, making it easier to use.
- There is a new option in SMIT to allow you to swap, in software, an active service or boot adapter to a standby adapter (`smit hacmp -> Cluster System Management -> Swap Adapter`).
- The location of HACMP log files can now be defined by the user by specifying an alternative path name via SMIT (`smit hacmp -> Cluster System Management -> Cluster Log Management`).
- Emulation of events for AIX error notification.
- Node-by-node Migration to HACMP ES is now included to assist customers who are planning to migrate from the HAS or CRM features to the ES or ESCRM features, respectively. Prior to this, conversion from HACMP to HACMP ES was the only option and required the cluster to be stopped.
- C-SPOC is again enhanced to take advantage of improvements to the LVM with the ability to make a disk known to all nodes and to create new shared or concurrent volume groups.

- An online cluster planning work sheet program is now available.
- The limitations detailed for HACMP ES Versions 4.2.1 and 4.2.2 have been addressed with the exception of the following:
 - VSM is not supported with a cluster greater than eight nodes.
 - Enhanced Scalability cannot execute in the same cluster as the HAS, CRM, or HANFS for AIX features.
 - The forced down function is not supported.
 - SOCC and slip networks are not supported.

A.4.8 HACMP and HACMP ES Version 4.4

HACMP Version 4.4 was introduced June 23, 2000 for RS/6000 uniprocessor and SMP servers and July 28, 2000 for RS/6000 SP systems. It provides the following enhancements to the previous release:

- Application Monitoring (ES and ESCRM only) provides Process Application monitoring and User-defined Application monitoring to determine the state of an application and to restart the application or fall the resource group over to another node.
- Tivoli Cluster Monitoring allows users to monitor the state of an HACMP cluster and its components on a Tivoli Desktop window.
- Cascading without Fallback permits specifying that the resource group not return to the original node when that node rejoins the cluster.
- Cluster Verification is enhanced to detect additional startup of failover problems.
- New Performance Tuning Parameters provide easier and more granular control.
- New documentation provides guidance for setting up, monitoring, and managing 24 x 7 clusters.
- Enhanced LVM Task Guide now provides a display of the physical location of each available disk and will create automatically a JFS log file.
- NFS Migration (HAS only) provides for migration between HACMP 4.3.1 HANFS feature and HACMP 4.4 HAS feature. HANFS is no longer included as a feature of HACMP.
- CSPOC is enhanced for specifying a user password and file system creation.

Appendix B. SCSI Cabling Examples

The following sections discuss various cabling examples using the following adapters:

- PCI Single-Ended Ultra SCSI Adapter
- PCI Differential-Ended Ultra SCSI Adapter

B.1 Cabling the PCI Single-Ended Ultra SCSI Adapter

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

- Ultra SCSI SE – Maximum transfer rate of 40 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-II Fast or Ultra SCSI devices is 6 meters (approximately 20 feet).
- The maximum supported cable length for configurations that include SCSI-II Fast (but not Ultra) is 3 meters (approximately 10 feet), with the exception of the 7027-HSC High Capacity Drawer, which can be attached with up to 6 meters of cable.
- To ensure optimum signal quality for Ultra SCSI transfers, attachment of multiple Ultra SCSI devices is only recommended for devices mounted inside the system unit. To ensure optimum signal quality for Ultra SCSI transfers, it is recommended that only Ultra SCSI devices be attached to a backplane that is driven by a PCI Single-Ended Ultra SCSI Adapter.

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

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

- Internal Ultra devices running at Ultra speeds:
 - Up to six Ultra devices attached to the internal port (dependent on internal configuration and cabling)
 - No external attachments allowed

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

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

B.1.1 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)

B.1.1.1 Adapter-to-First Device Cables

Table 173 describes the cables for the Single-Ended Ultra SCSI Adapter for connection to the first device.

Table 173. Single-Ended Ultra SCSI Adapter-to-First Device Cables

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

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

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

B.1.1.2 Device-to-Device Cables

Table 174 describes the device-to-device cables for single-ended applications.

Table 174. Device-to-Device Cables for Single-Ended Installations

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

Note

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

18.2.0.1 System-To-System Cables

Table 175 on page 637 provides a 16-bit connection between any two differential or single ended SCSI devices having 68-pin connectors. It can be

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

Table 175. 16-Bit SCSI-II System-to-System Cable

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

18.2.0.2 Terminators for Use with These Adapters

Table 176 describes the terminators for single-ended installations.

Table 176. Terminators for Single-Ended Installations

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

B.2 Cabling the PCI Differential-Ended Ultra SCSI Adapter

To understand the cabling for the PCI Differential Ultra SCSI Adapter (# 6207), first read 11.1.6, "General SCSI Cabling Considerations" on page 400, and then read the following for specific information.

B.2.1 SCSI Differential Cable Lengths Using This Adapter

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

B.2.1.1 Adapter-to-First Device Cables

Table 177 provides a list of adapter to first device cables.

Table 177. Differential-Ended Ultra SCSI Adapter-to-First Device Cables

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

When cables are ordered by feature code, the appropriate terminator is included with the order. When cables are ordered by part number, only the cable is included. For terminator part numbers, refer to the Table 179 on page 640.

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

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

B.2.1.2 Device-to-Device Cables

Table 178 provides a list of device-to-device cables.

Table 178. Device-to-Device Cables for Differential-Ended Installations

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

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

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

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

B.2.1.3 Terminators for Use With This Adapter

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

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

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

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

Table 179 provides a list of required terminators.

Table 179. Terminator for Description-Ended Installations

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

B.2.2 Cabling Examples for the PCI Differential Ultra SCSI Adapter

Figure 95 and Figure 96 on page 641 show how to cable the PCI Differential Ultra SCSI Adapter.

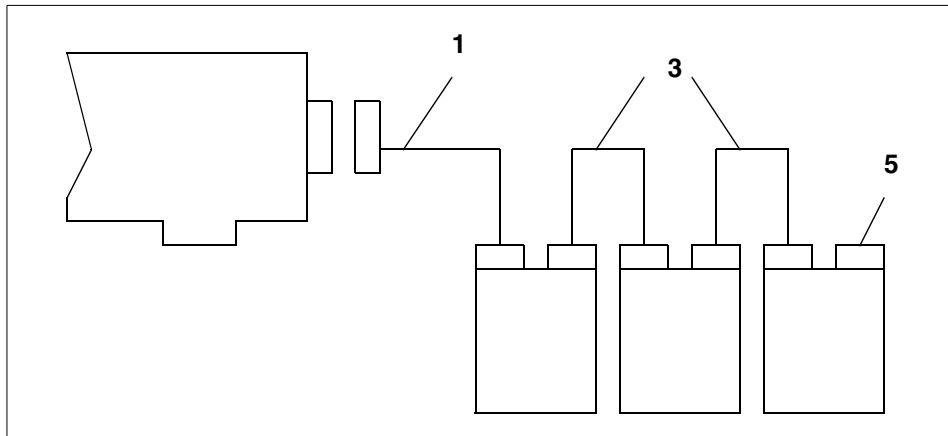


Figure 95. Differential External Narrow Bus




Figure 96. Differential External Wide Bus

Table 180 provides a list of available cables for differential adapters.

Table 180. Cabling for the PCI Differential Ultra SCSI Adapter

Item Number	Part Number	Length (meters)	Description
1	06H6037	1.0	Adapter-to-dual-connector device (narrow 8-bit)

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

Appendix C. RS/6000 Performance Data

The following tables contain performance data for all current RS/6000 models as well as for most older RS/6000 models. The data is organized into the following sections:

- SPEC and ROLTP benchmark data for uniprocessor systems
- SPEC and ROLTP benchmark data for SMP systems
- SPECweb96 benchmark data for selected systems
- TPC-C benchmark data for selected systems
- TPC-D benchmark data for selected systems
- LINPACK benchmark data for selected systems
- Notesbench Published Results

Appendix C.8, "Explanations of Performance Benchmarks" on page 653, gives definitions of the benchmarks listed in the following tables.

C.1 SPEC and ROLTP Performance Data - Single Processor Comparison

Table 181 and Table 182 list SPEC and ROLTP performance data for RS/6000 systems in a one processor configuration by model number.

* Indicates current machines

Table 181. SPEC and ROLTP Performance Data for Uniprocessor Systems

Model	L2 MB cache	MHz	CPU	SPEC int95	SPEC fp95	SPEC int base95	SPEC fp base95	SPEC int92	SPEC fp_92	Rel OLTP
220	0	33.3	POWER	0.6	0.9			20.4	29.1	0.3
230	0.13	33.3	POWER	0.9	1.2			28.5	39.9	0.4
250	0	66.7	601	1.82	2.32	1.69	2.32	62.6	76.0	1.0
250	0	80	601			2.03	2.58	78.8	90.4	
320	0	20	POWER	0.4	1.1			15.9	53.1	0.3
32H	0	25	POWER	0.6	1.7			21.5	45.3	0.4
340	0	33.3	POWER	0.7	2.2			29.7	64.5	0.7
34H	0	41.6	POWER	1.3	2.9			48.1	83.3	1.3
350	0	41.6	POWER	0.9	2.6			35.4	74.2	0.9
355	0	41.6	POWER	1.3	2.9			40.7	83.3	1.1
360	0	50	POWER	1.5	3.4			48.7	99.2	1.6
365	0	50	POWER	1.5	3.4			48.7	99.2	1.6
36T	0	50	POWER	1.5	3.4			48.7	99.2	1.6
370	0	62.5	POWER	1.9	4.2			70.3	121.1	1.7
375	0	62.5	POWER	1.9	4.2			70.3	121.1	1.7

Model	L2 MB cache	MHz	CPU	SPEC int95	SPEC fp95	SPEC int base95	SPEC fp base95	SPEC int92	SPEC fp_92	Rel OLTP
37T	0	62.5	POWER	1.9	4.2			70.3	121.1	1.7
380	0	59	POWER2	2.8	7.2			99.3	187.2	2.3
390	0	67	POWER2	3.14	7.50	3.00	7.20	109.7	202.1	2.6
390	0.5	67	POWER2					113.2	204.5	
390	1	67	POWER2	3.25	7.71	3.12	7.35	114.3	205.3	3.0
39H	0	67	POWER2					122.2	244.6	
39H	1	67		3.42	10.23	3.29	9.61	129.1	260.7	3.3
39H	2	67	POWER2	3.3	9.4			130.2	266.6	3.5
3AT	0	59	POWER2	2.84	7.20	2.71	6.81	99.3	187.2	2.3
3BT	0	67	POWER2	3.14	7.50	3.00	7.20	109.7	202.1	
3BT	1	67	POWER2	3.25	7.71	3.12	7.35	114.3	205.3	3.0
3CT	0	67	POWER2					122.2	244.6	
3CT	2	67	POWER2	3.42	10.23	3.29	9.61	130.2	266.6	3.5
397	0	160	P2SC	8.62	26.60	7.77	23.00			6.7
40P	0	66	601					64.2	68.3	
40P	0.25	66	601	1.92	2.4			76.0	77.2	
41W/T	0	80	601	2.03	2.58	2.03	2.58	78.8	90.4	1.4
41W/T	0.5	80	601	2.37	2.97			88.1	98.7	
42W/T	0	120	604	3.21	2.74	2.93	2.57	118.2	116.5	1.9
42W/T	0.5	120	604	4.01	3.53	3.75	3.37	150.2	146.5	2.4
43P-100	0.25	100	604	3.59	3.2	3.36	3.04	128.1	120.2	1.5
43P-120	0.5	120	604	4.24	3.41	4.01	3.23	157.9	139.2	1.9
43P-132	0.5	133	604	4.72	3.76	4.55	3.59	176.4	156.5	2.1
43P-140	0.5	166	604e	6.15	4.83	5.82	4.66			2.9
43P-140	1	200	604e	7.79	5.43	6.99	5.12			3.6
43P-140*	1	233	604e	9.24	5.75	8.29	5.48			3.9
43P-140*	1	332	604e	12.90	6.21	12.20	5.99			5.3
43P-150*	1	250	604e	11.1	8.78	10.9	8.28			4.0
43P-150*	1	375	604e	15.10	10.10	14.50	9.76			6.0
44P-170*	1	333	POWER3-II	19.8	35.6	18.6	34.0			10.4
44P-170*	4	400	POWER3-II	25.3	47.9	23.5	46.0			14.5
43P-240	0.5	166	604e	5.73	4.75					2.8
43P-240	1	233	604e	8.71	5.87	7.80	5.60			3.7
43P-260*	4	200	POWER3	12.5	27.0	11.5	25.1			10.5
44P-270*	4	375	POWER3-II	24.3	48.2	22.6	46.0			21.0
520	0	20	POWER	0.4	1.1				52.9	0.3
52H	0	25	POWER	0.6	1.7			22.5	49.3	0.4
530	0	25	POWER	0.5	1.6			20.1	72.5	0.5
53H	0	33.3	POWER	0.6	1.9			28.5	64.6	0.6
540	0	30	POWER	0.6	1.9			23.7	54.9	0.6
550	0	41.6	POWER	1.0	3.0			36.2	81.8	1.2
55L	0	41.6	POWER	1.3	2.9			40.7	83.3	1.1

Model	L2 MB cache	MHz	CPU	SPEC int95	SPEC fp95	SPEC int base95	SPEC fp base95	SPEC int92	SPEC fp_92	Rel OLTP
560	0	50	POWER	1.2	3.6			43.9	105.2	1.4
570	0	50	POWER	1.5	3.4			57.5	99.2	1.6
580	0	62.5	POWER	1.9	4.6			73.3	134.6	2.1
58H	0	55	POWER2	2.6	8.0			97.6	203.9	3.2
590	0	67	POWER2	3.33	10.38	3.19	9.69	121.6	259.7	3.9
591	0	77	POWER2	3.84	12.40	3.67	11.20	143.5	307.9	4.5
59H	1	67	POWER2	3.37	9.80	3.25	9.23	122.4	250.7	4.4
595	0	135	P2SC	6.17	17.60	5.90	15.40	240.0	470.0	5.8
730	0	25	POWER					20.1	72.5	
850	0.25	100	604					128.1	120.2	
850	0.5	120	604					157.9	139.2	
850	0.5	133	604					176.4	156.5	
930	0	25	POWER	0.5	1.6			20.1	72.5	0.5
950	0	41.6	POWER	1.0	3.0			38.0	88.3	1.2
970	0	50	POWER	1.3	3.3			47.8	101	1.4
97B	0	50	POWER	1.6	3.8			58.8	108.9	1.7
980	0	62.5	POWER	1.9	4.6			59.2	124.8	2.1
98B	0	62.5	POWER					73.3	134.6	
990	0	71.5	POWER2	3.4	11.2			131.0	279.0	4.2
B50*	1	375	604e	15.1	10.1	14.5	9.76			6.0
C10	0	80	601	2.03	2.58			78.8	90.4	1.4
C10	1	80	601	2.4	3.0	2.37	2.97	90.5	100.8	1.6
C20	0	120	604	3.2	2.7			118.2	116.5	1.7
C20	1	120	604	3.9	3.5	3.85	3.50	155.0	150.2	2.1
E20	0.5	100	604	3.67	3.13	3.43	3.06	139.6	131.6	2.5
E30	0.5	133	604	4.74	3.49	4.56	3.34			2.8
E30	1	166	604e	6.19	4.77	5.93	4.62			3.7
E30	0.5	233	604e	9.41	6.01	8.46	5.71			4.7
F30	0.5	133	604	4.74	3.49	4.56	3.34			2.8
F40	0.5	166	604e	5.73	4.75	5.31	4.60			2.8
F40	1.0	233	604e	8.71	5.34	7.80	5.12			3.7
F50*	0.25	166	604eX5	7.52	8.52	6.79	8.11			8.2
F50*	0.25	332	604eX5	14.40	12.60	14.00	12.10			10.0
F80*	2	450	RS64 III	21.0	25.4	18.7	24.8			23.0
F3L	0.5	233	604e	9.4	6.0					4.7
G30	0.5	75	601							
G40		112	604	3.7	3.1					2.6
G40		187	604e							
H10	0.5	166	604e	5.73	4.75	5.31	4.60			
H50*	0.25	332	604eX5	14.40	12.60	14.00	12.10			10.0
H70*	4	340	RS64 II	16.00	21.20	13.70	20.20			16.7
H80	2	450	RS64 III	21.0	25.4	18.7	24.8			23

Model	L2 MB cache	MHz	CPU	SPEC int95	SPEC fp95	SPEC int base95	SPEC fp base95	SPEC int92	SPEC fp_92	Rel OLTP
M20	0	33.3	POWER	0.6	0.9			17.3	29.1	0.3
M80*	4	500	RS64 III	24.1	29.1	20.7	28.5			
N40	0	50	601					41.7	51.0	
R10	0	50	POWER	1.5	3.4			57.5	99.2	1.6
R20	1	66.7	POWER2	3.37	9.80	3.25	9.23	122.4	250.7	4.4
R21	0	77	POWER2	3.84	12.40	3.67	11.20	143.5	307.9	4.5
R24	2	71.5	POWER2	3.47	10.20	3.32	9.47	134.1	273.8	4.9
SP T67	1	67	POWER2	3.25	7.71	3.12	7.35	114.3	205.3	
SP T67	2	67	POWER2	3.42	10.23	3.29	9.61	129.1	260.7	
SP W77	0	77	POWER2	3.84	12.40	3.67	11.20	143.5	307.9	4.5
SP T120	0	120	P2SC	5.61	16.60	5.36	14.60			5.8
SP W135	0	135	P2SC	6.17	17.60	5.90	15.40			5.8
SP T160*	0	160	P2SC	8.62	26.6	7.77	23.6			6.7
SP T/W*	.25	332	604e	14.4	12.6	14.0	12.1			6.7
SP T/W*	4	200	POWER3	12.5	27.0	11.5	25.1			10.5
SP H	4	222	POWER3	13.9	28.6	12.8	26.3			10.5
SP T/W	8	375	POWER3-II	24.4	50.96	22.6	47.1			10.5
SP H	8	375	POWER3-II	23.5	51.3	21.8	48.4			10.5
TP820	0.25	100	603					72.2	64.3	
TP850	0.25	100	603			2.53	2.07	102.3	99.1	
TP860	0.25	166	603EV	3.94	2.71	3.62	2.62			

Table 182. SPEC2000 for Uniprocessor Systems

Model	L2 MB cache	# of CPU	MHz	CPU	SPEC int 2000	SPEC fp 2000	SPEC int base2000	SPEC fp base2000
43P-150*	1	1	250	604e	105	90.8	99.4	90.8
43P-170*	1	1	333	POWER3-II	180	266	177	225
43P-170*	4	1	400	POWER3-II	249	344	239	295
44P-260*	4	1	200	POWER3				180
44P-270*	4	1	375	POWER3-II	251	360	242	309
F80*	2	1	450	RS64 III	234	210	225	205
H70*	4	1	340	POWER3-II			168	
H80*	2	1	450	RS64 III	234	210	225	205
M80*	4	1	500	RS64 III	275	250	264	243
SP T/W*	8	1	375	POWER3-II	260	382	248	330
SP H*	8	1	375	POWER3-II	252	337	229	322

C.2 SPEC and ROLTP Performance Data for SMP Systems

Table 183 and Table 184 list SPEC and ROLTP performance data for RS/6000 SMP systems by model number. * Indicates current machines

Table 183. SPEC and ROLTP Performance Data for SMP Systems

Model	L2 MB cache	# of CPU	MHz	CPU	SPEC int rate95	SPEC fp rate95	SPEC int rate base95	SPEC fp rate base95	SPEC int92	SPEC fp92	Rel OLTP
43P-240		1	233	604e	78.0	52.0	69.9	50.1			3.7
43P-240		2	233	604e	151.0	95.2	132.0	89.5			5.2
43P-260*	4	1	200	POWER3	112	243	104	225			10.5
43P-260*	4	2	200	POWER3	221	464	205	434			21.0
44P-270*	4	1	375	POWER3-II	218.0	434	203	414			21.0
44P-270*	4	2	375	POWER3-II	437.0	767	406	749			40.3
44P-270*	4	3	375	POWER3-II							54.2
44P-270*	4	4	375	POWER3-II	872	1146	811	1142			68.0
F40		2	166	604e	97.4	76.5	90.4	73.5			4.2
F40	1	1	233	604e	78.0	52.0	69.9	50.1			3.7
F40	1	2	233	604e	151	89.0	132	85.1			5.2
F50		1	166	604eX5	67.5	76.5	61.0	72.8			8.2
F50		2	166	604eX5	135	149	121	143			14.9
F50		3	166	604eX5							21.0
F50		4	166	604eX5	267	283	241	267			27.1
F50*	0.25	1	332	604eX5	128	116	124	109			10.0
F50*	0.25	2	332	604eX5	255	218	245	206			17.9
F50*	0.25	3	332	604eX5	380	310	365	292			25.2
F50*	0.25	4	332	604eX5	485	364	389	364			32.8
F80*	2	1	450	RS64 III	189	229	168	223			23.0
F80*	4	2	450	RS64 III							50.0
F80*	4	4	450	RS64 III	783	837	674	821			87.7
F80*	4	6	500	RS64 III	1298	1153	1118	1135			111.9
G30		2	75	601	40.4	44.1	36.9	43.2	3840	4040	3.3
G30		4	75	601	77.1	81.5	73.1	78.8	7580	7280	5.9
G40		2	112	604	66.5	53.3	60.6	50.7			4.8
G40		4	112	604	122		110				8.8
G40		2	187	604e	129	87.1	113	80.7			6.6
G40		4	187	604e	254	154	217	149			
H10		1	233	604e	78.0	52.0	69.9	50.1			3.7
H10		2	233	604e	151	95.2	132	89.5			5.2
H50*	0.25	1	332	604eX5	128	116	124	109			10.0
H50*	0.25	2	332	604eX5	255	218	245	206			17.9
H50*	0.25	3	332	604eX5	380	310	365	292			25.2
H50*	0.25	4	332	604eX5	501	389	485	364			32.8
H70*	4	1	340	RS64 II	144	191	124	182			16.7

Model	L2 MB cache	# of CPU	MHz	CPU	SPEC int rate95	SPEC fp rate95	SPEC int rate base95	SPEC fp rate base95	SPEC int92	SPEC fp92	Rel OLTP
H70*	4	2	340	RS64 II	287	370	247	354			31.9
H70*	4	3	340	RS64 II	430	534	370	512			44.5
H70*	4	4	340	RS64 II	573	674	492	645			57.1
H80*	2	1	450	RS64 III	189	229	168	223			23.0
H80*	4	2	450	RS64 III							50.0
H80*	4	4	450	RS64 III	783	837	674	821			87.7
H80*	4	6	500	RS64 III	1298	1153	1118	1135			111.9
J30		2	75	601	42.3	47.4	39.7	46.4	4282	4492	3.7
J30		4	75	601	83.0	91.3	78.2	89.3	8430	8689	6.6
J30		6	75	601	125	134	117	133	12008	12644	9.4
J30		8	75	601	162	172	153	165	16200	16324	12.1
J40	1	2	112	604	71.9	57.3	64.9	53.4			5.8
J40	1	4	112	604	138	107	129	102			10.0
J40	1	6	112	604	205	159	195	154			14.5
J40	1	8	112	604	258	200	244	189			19.2
J50		2	200	604e	137	92.5	121	90.3			9.3
J50		4	200	604e	268	183	244	176			17.0
J50		6	200	604e	396	261	343	248			23.8
J50		8	200	604e	509	332	445	320			30.6
M80*	4	2	500	RS64 III							65.0
M80*	4	4	500	RS64 III							115.3
M80*	4	5	500	RS64 III							169.3
M80*	4	8	500	RS64 III	1728	1958	1489	1910			222.5
R30		2	75	601	42.3	47.4	39.7	46.4	4282	4492	3.7
R30		4	75	601	83.0	91.3	78.2	89.3	8430	8689	6.6
R30		6	75	601	125	134	117	133	12008	12644	9.4
R30		8	75	601	162	172	153	165	16200	16324	12.1
R40	1	2	112	604	71.9	57.3	64.9	53.4			5.8
R40	1	4	112	604	138	107	129	102			10.0
R40	1	6	112	604	205	159	195	154			14.5
R40	1	8	112	604	258	200	244	189			19.2
R50		2	200	604e	137	92.5	121	90.3			9.3
R50		4	200	604e	268	183	244	176			17.0
R50		6	200	604e	396	261	343	248			23.8
R50		8	200	604e	509	332	445	320			30.6

Model	L2 MB cache	# of CPU	MHz	CPU	SPEC int rate95	SPEC fp rate95	SPEC int rate base95	SPEC fp rate base95	SPEC int92	SPEC fp92	Rel OLTP
S70	4	4	125	RS64	SPEC data is available for the S70/S7A through the RS/6000 Special Bid Web page http://rs6000bid.austin.ibm.com . Bear in mind that the S70/S7A family is positioned for commercial workloads, such as OLTP, ERP, BI, e-business, and server consolidation. Appropriate industry standard benchmark and application performance benchmark data should be used instead to assess performance in these particular areas.						24.6
S70	4	8	125	RS64							46.3
S70	4	12	125	RS64							62.2
S7A*	8	4	262	RS64 II							52.7
S7A*	8	8	262	RS64 II							98.7
S7A*	8	12	262	RS64 II							136.7
S80	8	6	450	RS64 III							139
S80	8	12	450	RS64 III							265
S80	8	18	450	RS64 !!!							369.1
S80	8	24	450	RS64 III							452.7
SP Hi	1	2	112	604	71.9	57.3	64.9	53.4			5.8
SP Hi	1	4	112	604	138	107	129	102			10.0
SP Hi	1	6	112	604	205	159	195	154			14.5
SP Hi	1	8	112	604	258	200	244	189			19.2
SP H2		2	200	604e	137	92.5	121	90.3			9.3
SP H2		4	200	604e	268	183	244	176			17.0
SP H2		6	200	604e	396	261	343	248			23.8
SP H2		8	200	604e	509	332	445	320			30.6
SP T/W	0.25	2	332	604eX5	255	218	245	206			17.9
SP T/W	0.25	4	332	604eX5	501	389	485	364			32.8
SP T/W	4	1	200	POWER3	112	243	104	225			10.5
SP T/W	4	2	200	POWER3	221	464	205	434			21.0
SP H	4	2	222	POWER3	249	503	229	461			23.0
SP H	4	4	222	POWER3	493	999	450	910			43.3
SP H	4	6	222	POWER3	724	1464	661	1329			64.0
SP H	4	8	222	POWER3	966	1877	908	1760			81.3
SP T/W	8	2	375	POWER3-II	438	844	407	804			44.0
SP T	8	4	375	POWER3-II	875	1382	812	1359			67.7
SP W	8	4	375	POWER3-II	875	1382	812	1359			80.0
SP W-8	8	4	375	POWER3-II	6014	10605	5616	10395			
SP W-16	8	4	375	POWER3-II	10366	20251	10001	20220			
SP W-32	8	4	375	POWER3-II	17920	38073	17239	38073			
SP H*	8	4	375	POWER3-II	845	1739	786	1670			81.7
SP H*	8	8	375	POWER3-II	1684	3418	1569	3290			163.7
SP H*	8	12	375	POWER3-II	2523	4985	2345	4832			242.3
SP H*	8	16	375	POWER3-II	3352	6353	3121	6202			319.3

Table 184. SPEC2000 for SMP Systems

Model	L2 MB cache	# of CPU	MHz	CPU	SPEC int rate 2000	SPEC fp rate 2000	SPEC int rate base2000	SPEC fp rate base2000
44P-270*	4	2	375	POWER3-II	5.75	7.56	5.54	6.53
44P-270*	4	4	375	POWER3-II	11.2	11.8	10.8	10.5
F80*	4	1	450	RS64 III	2.72	2.44	2.62	2.38
F80*	4	2	450	RS64 III	11.3	9.45	10.8	9.26
F80*	4	4	500	RS64 III	18.0	13.2	17.3	13.0
H80*	4	1	450	RS64 III	2.72	2.44	2.62	2.38
H80*	4	2	450	RS64 III	11.3	9.45	10.8	9.26
H80*	4	4	500	RS64 III	18.0	13.2	17.3	13.0
M80*	4	8	500	RS64 III	25.1	21.1	24.0	20.6
SP H*	8	4	375	POWER3+	11.6	14.5	10.6	14.1
SP H*	8	8	375	POWER3+	23.1	28.0	21.0	27.0
SP H*	8	12	375	POWER3+	34.6	41.1	31.4	39.0
SP H*	8	16	375	POWER3+	46.0	51.7	41.7	49.7

C.3 SPECweb96 and SPECweb99 Benchmarks

Table 185 lists the SPECweb96 and SPECweb99 benchmarks for current RS/6000 systems.

Table 185. Web Benchmarks for Selected RS/6000 Systems

Model	L2 cache	# of CPU	MHz	CPU	SPECweb96 (ops/sec.)	SPECweb99 (con/sec.)
43P-140	0.5	1	166	604e	459	
43P-260	4	2	200	POWER3	4597	440
44P-170				POWER3-II		460
44P-270				POWER3-II		1359
F50	0.25	4	166	604eX5	2148	
F50/H50	0.25	4	332	604eX5	2755	
H70	4	2	340	RS64 II	6958	
H70	4	4	340	RS64 II	11,774	
H80	4	6	500	RS64 III		2200
M80	4	8	500	RS64 III		3216
S70	4	12	125	RS64	4075	
S70u	8	12	262	RS64 II	19,264	
S7A	8	12	262	RS64 II	20,200	
S80	8	12	450	RS64 III	40,161	
SP T/W	.25	4	332	604eX5	6716	
SP T/W	4	2	200	POWER3	4597	

C.4 TPC-C Benchmark Data

Table 186 provides the tpmC and \$/tpmC benchmark for selected RS/6000 systems. * Indicates current machines

Table 186. TPC-C Benchmark Results

Model	Processor	MHz	L2 cache	tpmC	\$/tpmC	Database	AIX Vers.
C10	601	80	1	485.88		Sybase V10	3.25
E20	604	100	0.5	735.27		Sybase V10	4.14
R24	POWER2	71	2	1470.06		DB2 V2.1	3.2.5
R40	604/8	112	1	5774.07		Sybase V11	4.1.4
R50	604/8	200	2	9165.13	98.83	Sybase V11	4.2.1
J40	604/8	112	1	5774.07		Sybase V11	4.1.4
J50	604e/8	200	2	9165.13		Sybase V11.5	4.2.1
F50*	604e/4	166	0.25	8142.40	62.71	Sybase V11.5	4.2.1
F50*	604e/4	332	0.25	9853.13	64.22	Sybase V11.5	4.2.1
F80*	RS64 III/6	500	4	33571.39	58.94	Oracle V816	4.3.3
H70*	RS64 II/4	340	4	17133.73	78.41	Oracle V815	4.3.2
M80*	RS64 III/8	500	4	66750.27	45.46	Oracle V817	4.3.3
S70	RS64/12	125	4	18666.73	108.62	Oracle 8.0	4.3.0
S70	RS64 II/12	262	8	34139.63	88.09	Oracle 8.0	4.3.1
S7A*	RS64 II/12	262	8	110434.10	122.44	Oracle OPS	4.3.2
S80	RS64 III/24	450	8	135,815.7	52.70	Oracle8i 8.1.6	4.3.3

C.5 TPC-D Published Results (Obsolete as of 8/15/99)

All Results are withdrawn as obsolete benchmarks (replacement is TPC-H, TPC-R). Table 187 and Table 188 lists the TPC-D results for tested RS/6000 systems.

Table 187. TPC-D Published Results

Model	Process.	No. of Nodes	MHz	Power QppD@ 100GB	Through-put/QthD @100GB	Price/ Perf\$/ QphD@ 100GB	Streams	RDBMS	AIX Vers.	Avail. Date
F50	604e/4	1	332	1168.1	499.1	429	1	DB2 UDB	4.2.1	03/31/98
H50	604e/4	1	332	1168.1	499.1	446	1	DB2 UDB	4.2.1	03/31/98
F50	604e/4	1	332	1245.5	562.0	433	5	DB2 UDB	4.2.1	10/31/98
H50	604e/4	1	332	1245.5	562.0	449	5	DB2 UDB	4.2.1	10/31/98
S7A	RS64 II/12	1	262	4226.5	1092.6	589	5	DB2 UDB	4.3.2	12/31/98

Table 188. TPC-D Published Results for SP Nodes

Model	Process.	No. of Nodes	MHz	Power QppD@ 300GB	Through-put/QthD @300GB	Price/ Perf\$/ QphD@ 300GB	Streams	RDBMS	AIX Vers.	Avail. Date
SP	604e/4	24	332	10469.6	6166.5	721	16	DB2 UDB	4.2.1	10/31/98

Model	Process.	No. of Nodes	MHz	Power QppD@ 1TB	Through-put/QthD @1TB	Price/ Perf\$/ QphD@ 1TB	Streams	RDBMS	AIX Vers.	Avail. Date
SP	604e/4	48	332	19137.5	10661.5	797		DB2 UDB	4.2.1	10/31/98

C.6 LINPACK Benchmark Data

Table 189 lists LINPACK performance data for current RS/6000 uniprocessor systems.

Table 189. LINPACK Performance Data for Uniprocessor Systems

Model	CPU	MHz	L1 Cache (KB)	L2 Cache (MB)	LINPACK DP	LINPACK SP	LINPACK TPP
43P-140u	604e	233	32/32	1.0	22.6	77.1	156.2
43P-140n	604e	233	32/32	1.0	56.0	113.3	156.2
43P-140	604e	332	32/32	1.0	59.9	123.6	179.7
43P-150	604e	375	32/32	1.0	64.8	151.2	255.7
44P-170	POWER3-II	333	32/64	1.0	363		833
44P-170	POWER3-II	400	32/64	4.0	461		1052
43P-260	POWER3	200	32/64	4.0	236.5	248.1	630.0

Model	CPU	MHz	L1 Cache (KB)	L2 Cache (MB)	LINPACK DP	LINPACK SP	LINPACK TPP
44P-270	POWER3-II	375	32/64	4.0	425		1045
B50	604e	375	32/32	1.0	64.8	151.2	255.7
F40	604e/1	233	32/32	1.0	48.5	109.1	145.6
F50	604e/1	166	32/32	0.2	70.2	98.9	166.4
F50	604e/1	332	32/32	0.2	115.7	158.5	273.4
H50	604e/1	332	32/32	0.2	115.7	158.5	273.4
H70	RS64 II	340	64/64	4.0	187.6	124.0	498.3

C.7 Notesbench Published Results

Table 190 lists Notesbench Published Results for tested RS/6000 systems.

Table 190. Notesbench Published Results

Model	CPU	# of CPU	MHz	L1 Cache (KB)	L2 Cache (MB)	Users	TPM	Resp Time	\$/User	Version
43P-140n	604e	1	233	32/32	1.0	1450	1917	0.484	11.97	4.52
F50	604e	1	332	32/32	0.2	6000	7947	0.406	14.87	4.53b
F50	604e	1	332	32/32	0.2	6400	8919	0.292	16.15	4.6
H70	340	4	340	64/64	4.0	15372	11000		19.65	4.6
S70	RS64 II	12	262	64/64	8.0	28800	40075	0.213	21.32	4.6
S80	RS64 III	24	450	128/128	8.0	57600	71904		27.51	5.0
M80	RS64 III	8	500	128/128	4.0	28032	38235	1.424	23.91	5.0

C.8 Explanations of Performance Benchmarks

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

SPEC

<http://www.specbench.org>

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX Version 4.2 or Version 4.3. IBM C for AIX Version 3.6.4.0 and XL Fortran Version 6.1.0.0 and Version 7.1.0.0 were the compilers used in the benchmark tests. The preprocessors used in the benchmark tests include KAP 3.2 for Fortran and KAP/C 1.4.2 from Kuck and Associates and VAST-2 Version 4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors.

C.8.1 SPEC Benchmark Definition

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

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

SPEC95 contains two suites of benchmarks:

- CINT95** Measures and compares computational-intensive integer performance
- CFP95** Measures and compares computational-intensive floating point performance

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

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 outdated VAX 11/780 to a SPARCstation 10/40 with 64 MB memory, but without Level 2 cache. This is more

representative of the types of systems being sold today, but it is also a machine that will beat few, if any, of the machines being benchmarked.

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

The SPEC base metric (for example, SPECint_base95) is required for all reported results and has 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.

SPEC CPU2000 comprises two sets (or suites) of benchmarks: CINT2000 for measuring compute-intensive integer performance, and CFP2000 for compute-intensive floating point performance. The two suites measure the performance of a computer's processor, memory architecture and compiler. Improvements to the new suites include longer run times and larger problems for benchmarks, more application diversity, greater ease of use, and standard development platforms that will allow SPEC to produce additional releases for other operating systems.

SPEC CPU2000 benchmarks were selected from application and source codes running across multiple platforms. Seventeen (17) of the 26 benchmarks resulted from SPEC's search campaign that offered the general public rewards up to \$5,000 for application codes and datasets. Each benchmark was tested on different platforms to determine if it was portable, relevant, and suitable for the final SPEC CPU2000 suite.

SPEC CPU2000 is the fourth major version of the SPEC CPU benchmark suites, which, in 1989, became the first widely accepted standard for comparing compute-intensive performance across various architectures. The new release replaces SPEC CPU95, which will be gradually phased out by July 2000, when SPEC will stop publishing CPU95 results. Performance

results from CPU2000 cannot be compared to those from CPU95 since new benchmarks have been added and existing ones changed.

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

SPECint95	SPEC component-level benchmark that measures integer performance. Result is the geometric mean of eight tests that comprise the CINT95 benchmark suite. All of these are written in C language.
SPECint_base95	The result of the same tests in CINT95 with a maximum of four compiler flags that must be used in all eight tests.
SPECint_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95).
SPECint_base_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95) with the restrictive compiler options.
SPECfp95	SPEC component-level benchmark that measures floating point performance. The result is the geometric mean of ten tests that comprise the CFP95 benchmark suite. All of these are written in Fortran.
SPECfp_base95	Result of the same tests in CFP95 with a maximum of four compiler flags that must be used in all ten tests.
SPECfp_rate95	Geometric average of the ten SPEC rates from SPEC floating point tests (CFP95).
SPECfp_base_rate95	Geometric average of the ten SPEC rates from the SPEC floating-point tests (CFP95) with the restrictive compiler options.
SPECint2000	New SPEC component-level benchmark that measures integer performance. The result is the geometric mean of twelve tests that comprise the CINT2000 benchmark suite. All of these are written in C language except for one which is in C++.

SPECint_base2000	The result of the same tests in CINT2000 with a maximum of four compiler options that must be used in all twelve tests.
SPECint_rate2000	Geometric average of the twelve SPEC rates from the SPEC integer tests (CINT2000).
SPECint_base_rate2000	The result of the same tests as CINT2000 with a maximum of four compiler options that must be used in all twelve tests.
SPECfp2000	New SPEC component-level benchmark that measures floating point performance. The result is the geometric mean of fourteen tests, all written in Fortran and C languages, that are included in the CFP2000 benchmark suite.
SPECfp_base2000	Result of the same tests in CFP2000 with a maximum of four compiler options that must be used in all fourteen tests.
SPECfp_rate2000	Geometric average of the fourteen SPEC rates from SPEC floating point tests (CFP2000).
SPECfp_base_rate2000	The result of the same tests as CFP2000 with a maximum of four compiler options that must be used in all fourteen tests.
SPECweb96	Maximum number of Hypertext Transfer Protocol (HTTP) operations per second achieved on the SPECweb96 benchmark without significant degradation of response time. The Web server software is ZEUS v1.1 from Zeus Technology, Ltd. or IBM HTTP Server 1.3.4.
SPECweb99	The number of conforming, simultaneous connections the Web server can support using a predefined workload. The SPECweb99 test harness emulates clients sending the HTTP requests in the workload over slow Internet connections to the Web server. The Web server software is Zeus from Zeus Technology Ltd.

C.8.2 Relative Online Transaction Processing (ROLTP) Definition

Relative online transaction processing (ROLTP) is an estimate of commercial processing performance derived from an IBM analytical model. The model simulates some of the system's operations, such as CPU, cache, and

memory. However, the model does not simulate disk or network I/O operations. Although general database and operating systems parameters are used, the model does not reflect specific databases or AIX version or releases. Unless otherwise indicated for a system, the model assumes the use of 32-bit applications. ROLTP is estimated only at the time the system is introduced unless otherwise indicated for a system. An IBM RS/6000 Model 250 is the baseline reference system and has a value of 1.0.

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

C.8.3 Transaction Processing Council (TPC) Benchmark Definitions

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

- tpmC** TPC Benchmark C throughput measured as the average number of transactions processed per minute during a valid TPC-C configuration run of at least twenty minutes.
- \$/tpmC** 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 such estimated total cost by the tpmC for the system.
- QppD** 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** Throughput metric of TPC-D and is a classical throughput measure characterizing the ability of the system to support a multiuser 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 different 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** 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.

C.8.4 LINPACK Benchmark Definition

LINPACK is a collection of Fortran subroutines that analyze and solve linear equations and linear least-squares problems. The package solves linear systems whose matrices are general, banded, symmetric indefinite, symmetric positive definite, triangular, and tridiagonal square. In addition, the package computes the QR and singular value decompositions of rectangular matrices and applies them to least-squares problems. LINPACK uses column-oriented algorithms to increase efficiency by preserving locality of reference.

LINPACK was designed for supercomputers in use in the 1970s and early 1980s. LINPACK has been largely superseded by LAPACK, which is designed to run efficiently on shared-memory, vector supercomputers.

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

LINPACK SP	Single precision, n=100 results with AIX XL Fortran compiler with optimization. Units are megaflops (MFLOPS).
LINPACK DP	Double precision, n=100 results with AIX XL Fortran compiler with optimization. Units are megaflops (MFLOPS).
LINPACK TPP	Toward Peak Performance, n=1000 results with AIX XL Fortran compiler with optimization. Units are megaflops (MFLOPS). ESSL Version 3.1.1 was used in this test.

C.8.5 Notesbench

Notesbench is the driver program to test various aspects of Lotus Notes. It is designed to execute the commands in customized workload scripts, simulating Notes client actions. Notesbench tests include "mail only" and "mail and database". All IBM results have been published with the "mail only" workload.

The following NotesBench information provides consistent workload and performance information. These metrics are used to determine server capacity and performance.

Total Users	The number of active users supported in the workload, each producing approximately one transaction/minute.
--------------------	--

TPM	A throughput performance metric measured in transactions per minute. (NotesMark)
Average Response Time	The average time for a transaction to be completed for an average user action.
\$/User	The total cost of the hardware and the software, including discounts quoted by a supplier.

Appendix D. Adapter Placement Guide

This section summarizes the adapter placement rules for the following RS/6000 models:

- 7043 Model 140
- 7043 Model 150
- 7044 Model 170
- 7043 Model 260
- 7044 Model 270
- 7046 Model B50
- 7025 Model F50
- 7025 Model F80
- 7026 Model H50
- 7026 Model H70
- 7026 Model H80
- 7026 Model M80
- 7017 Models S70 and S70 Advanced
- 7017 Model S80

For an introduction to the PCI architecture, see 1.2.1.1, “PCI Slots” on page 3.

D.1 32-Bit Versus 64-Bit PCI Slots

A variable that effects slot placement and performance is the choice of 32-bit versus 64-bit PCI slots. 64-bit slots were designed for higher speed adapters because they can transfer 64 bits of data for each data transfer phase.

32-bit adapters can typically be installed in 64-bit PCI slots; however, 32-bit adapters still operate in 32-bit mode and have no performance advantage in a 64-bit slot. Likewise, most 64-bit adapters can be installed in a 32-bit PCI slot, but, when this is done, the 64-bit adapter operates in 32-bit mode, and the adapter performance is reduced.

D.2 33 MHz Versus 50 MHz 64-Bit PCI Slots

Some systems (for example, 7025 Model F50 and 7026 Model H50) offer 50 MHz capability on 64-bit slots. Adapters capable of functioning at 50 MHz may take advantage of this. If a 33 MHz adapter is plugged into a 50 MHz 64-bit slot, the slot switches to 33 MHz and also switches the remaining slots on this PCI bus to 33 MHz.

The following adapters run at 50 MHz when placed on a 50 MHz PCI bus or 33 MHz on a 33 MHz PCI bus. If a 33 MHz adapter is placed on the same 50 MHz PCI bus with any of these adapters, the bus is forced to run in 33 MHz mode, which reduces the performance of these adapters:

- Gigabit Ethernet PCI
- POWER GXT250P
- POWER GXT255P
- POWER GXT300P
- POWER GXT2000P
- POWER GXT3000P

D.3 Connectivity Versus Performance Overview

There are trade-offs when configuring a system regarding how many adapters of each type may be physically supported (installed and operational) versus how many can be supported for maximum performance. The following paragraphs provide an overview of these considerations and how they are documented in later sections of this appendix.

Connectivity limits define how many adapters of a specified type can be physically plugged into a system. This limit defines how many adapters the software and hardware can support. There may be specific guidelines as to the placement of these adapters as well. Connectivity limits define the maximum number of adapters for connecting to networks or disks. In many cases, the duty-cycle of a disk or network is low, and additional adapters are needed to obtain the physical connection to all resources. In these cases, connectivity limits should be followed.

Performance limits are suggested guidelines established to help determine how many adapters can be running concurrently with each adapter providing good performance. Another way to view this is that as you add adapters, with each adapter performing at close to its rated speed, each added adapter would continue to provide an incremental performance increase. Once the

performance limit is reached, adding more adapters does not provide an increase in I/O throughput.

A number of factors can determine the performance limit. It might be the bus speed, memory speed, adapter design, or processor speed. Quite often, the system processor speed may be a limiting factor on how many adapters of a given type can be supported and operated at maximum performance. Once a system is close to 90 percent of system processor utilization, adding more adapters only provides a minor increase in throughput.

Due to the wide variety of workloads, performance limits in this book are only guidelines. The guidelines are based on I/O streaming of large reads or writes to disk or to a network. They are not based on small I/Os, which are more transaction rate limited. Small I/O workloads probably utilize more system processor capacity and result in fewer adapters being supported for maximum performance.

These guidelines are also based on the maximum number of processors supported for multiprocessor systems. If your system is running less than the maximum number of processors supported, then, typically, you have to reduce the maximum number of adapters by the same ratio. For example, if a system with a maximum of twelve processors can support twelve ATM adapters for maximum performance, then the same system with eight processors can only support eight ATM adapters for maximum performance.

If the system will be using disk and communication adapters concurrently, for maximum performance a more conservative estimate of the number of adapters supported should be used.

If your configuration is close to the performance limits, then extra care and investigation should be done to ensure that the system type or configuration provides the desired performance. In these cases, you may need to contact your marketing support personnel for more detailed information.

D.4 Other Restrictions

Some adapters must be installed in specific PCI slots in various systems. This may be due to physical size limits, I/O address considerations, thermal limitations, and other factors. The *PCI Adapter Placement Reference*, SA38-0538 lists restrictions and slot placement guidelines for PCI adapters in various system units.

D.5 7043 Model 140 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 97. Model 140 System Unit Rear View with Numbered Slots

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 191 to determine where to install a single adapter in your system unit.

If you are running AIX, refer to Table 192 on page 666 to install two or more adapters in your 7043 Model 140. Windows NT (PowerPC Edition) does not support more than one graphics accelerator in a system unit.

Use Figure 97 to identify the slot locations described in the following tables.

Table 191. Model 140 Single Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)	NT
2962	2-Port Multiprotocol PCI	Slots 3P, 4P, 5P	
2947	IBM ARTIC960Hx 4-Port Selectable PCI	Slots 2P, 3P, 4P	
2948	IBM ARTIC960Hx 4-Port T1/E1 PCI	Slots 2P, 3P, 4P	
2949	IBM ARTIC960Hx DSP Resource PCI	Slots 2P, 3P, 4P	
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Slots 2P, 3P, 4P (see note 1)	
6309	Digital Trunk Quad PCI	Slots 2P, 3P, 4P (Max. 2 per system)	
2839	POWER GXT110P	Slots 1P-2P	Slots 1P-2P
2838	POWER GXT120P	Slots 1P-5P	
2830	POWER GXT130P	Slots 1P-5P	

Feature Code	Adapter	Slot (AIX)	NT
2851 or 2852	POWER GXT250P or POWER GXT255P	Slots 1P-5P	Slots 1P-2P
2841	POWER GXT300P	Slots 1P-5P	
2854 or 2855	POWER GXT500P or POWER GXT550P	Slot 2P	
2853	POWER GXT800P	Slot 2P (see note 2)	
2856	POWER GXT1000 Attachment Adapter	Slots 1P-2P	
2823	POWER GXT2000P	Slots 1P, 2P, 3P, 4P, 5P	
6215	PCI SSA Multi-Initiator/RAID EL	Slot 2P	
6218	SSA 4-Port RAID	Slot 2P	
6206	Single-Ended Ultra SCSI	Slot 5P	
2639	Ultimedia Video Capture	Slots 1P, 2P, 3P, 4P	
2963	Turboways 155 PCI UTP ATM	See note 2	
2988	Turboways 155 PCI MMF ATM	See note 2	
2968	10/100 Mbps Ethernet	See note 3	
4951	4-Port 10/100 Base-TX Ethernet	Slots 1P, 2P (Max. 2 per system, see note 5)	

Notes:

1. Do not install more than a combination of 3 IBM ARTIC960RxD Quad Digital Trunk PCIs and IBM ARTIC960RxF PCI Adapters in this system.
2. The POWER GXT800P takes the physical space of three adapters in the 7043 Model 140.
3. For optimum system performance, when using a 155 TURBOWAYS ATM PCI MMF adapter or 155 TURBOWAYS ATM PCI UTP adapter in LAN emulation (LANE), a maximum of one adapter per system is recommended.
4. A maximum of four 10/100 Mbps Ethernet PCI Adapters can be used in the 7043 Model 140. For optimum system performance, one 10/100 Mbps Ethernet PCI Adapter operating in 100 Mbps mode is recommended.

5. For optimum system performance, use only one 100 Mbps port per system.

Use Table 192 to identify adapter location combination options for these adapters in your 7043 Model 140.

Do not install more than four graphics cards in your 7043 Model 140.

Table 192. Model 140 Multiple Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)	Alternate Slot (AIX)
2839 with a 2838 or 2830 or 2851 or 2852	GXT110P with a GXT120P or GXT130P or GXT250P or GXT255P	Slot 1P Slot 2P	Slot 2P Slot 1P
2839 with a 2854 or 2855	GXT110P with a GXT500P or GXT550P	Slot 1P Slot 2P	
2839 with a 2853	GXT110P with a GXT800P (see note 1)	Slot 1P Slot 2P	
2839 with a 2856	GXT110P with a GXT1000 Attachment Adapter	Slot 1P Slot 2P	Slot 2P Slot 1P
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2838 or 2830 or 2851 or 2852 or 2841 or 2823	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P	Slot 1P-5P Slot 1P-5P	
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2854 or 2855	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT500P or GXT550P	Slot 1P Slot 2P	
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2853	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT800P (see note 1)	Slot 1P Slot 2P	
2639 with a 2853	Ultimedia Video Capture with a GXT800P (see note 1)	Slot 1P Slot 2P	
2838 or 2830 or 2851 or 2852 or 2841 with a 2856	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P with a GXT1000 Attachment Adapter	Slot 1P Slot 2P	Slot 2P Slot 1P

Feature Code	Adapter	Slot (AIX)	Alternate Slot (AIX)
2639 with a 2823	Ultimedia Video Capture with a GXT2000P	Slot 1P-4P Slot 1P-5P	
2823 with a 2853	GXT 2000P with a GXT800P (see note 1)	Slot 1P Slot 2P	
2639 with a 2841	Ultimedia Video Capture with a GXT300P	Slot 1P-4P Slot 1P-5P	

Notes:

1. The POWER GXT800P takes the physical space of three adapters in the 7043 Model 140.
2. A maximum of two 10/100 Mbps Ethernet PCI Adapters can be used in the 7043 Model 140. For optimum system performance, one 10/100 Mbps Ethernet PCI Adapter operating in 100 Mbps mode is recommended.

D.6 7043 Model 150 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 98. Model 150 System Unit Rear View with Numbered Slots

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 193 on page 668 to determine where to install a single adapter in your system unit and Table 194 on page 670 to determine where to install multiple adapters in your system unit.

Use the rear view diagram in Figure 98 on page 667 to identify slot locations described in the following tables.

Table 193. Model 150 Single Adapter Placement Guide

Feature Code	Adapter	Usage	System Maximum
2639	Ultimedia Video Capture	Slots 1P-5P	1
2838	POWER GXT120P	Slots 1P-5P	4
2830	POWER GXT130P	Slots 1P-5P	4
2851	POWER GXT250P	Slots 1P-5P	4
2852	POWER GXT255P	Slots 1P-5P	4
2841	POWER GXT300P ¹	Slots 1P-5P	4
2845	POWER GXT550P	Slots 2P, 3P	1
2823	POWER GXT2000P	Slots 1P-5P	4
2825	POWER GXT3000P ²	Slot 3P	1
2741	SysKonnnect SK-NET FDDI-LP SAS PCI ³	Slots 1P-5P	2
2742	SysKonnnect SK-NET FDDI-LP DAS PCI ³	Slots 1P-5P	2
2743	SysKonnnect SK-NET FDDI-UP SAS PCI ³	Slots 1P-5P	2
6206	PCI Single-Ended Ultra SCSI	Slots 1P-5P	2
6207	PCI Differential Ultra SCSI	Slots 1P-5P	2
2494	PCI 3-Channel Ultra SCSI RAID ¹⁰	Slots 1P-5P	2
6215	PCI SSA Multi-Initiator/RAID EL	Slots 2P, 3P	1 ⁷
6225	Advanced Serial RAID	Slots 2P, 3P, 4P	1 ⁷
6230	Advanced Serial RAID Plus	Slots 2P, 3P, 4P	1 ⁷
4959	Token ring PCI	Slots 1P-5P	4
2968	10/100 Mbps Ethernet PCI ⁴	Slots 1P-5P	4
4951	4-Port 10/100 Base-TX Ethernet PCI ¹¹	Slots 2P or 3P	2
2943	8-Port Asynchronous EIZ-232E/RS-422A PCI	Slots 1P-5P	2

Feature Code	Adapter	Usage	System Maximum
2944	128-Port Async Controller PCI	Slots 1P-5P	2
2962	2-Port Multiprotocol PCI	Slots 5P, 4P, 1P	2
2963	155 Turboways ATM PCI UTP	Slots 1P-5P	2 ⁵
2988	155 Turboways ATM PCI MMF	Slots 1P-5P	2 ⁵
2708	ISDN Basic Rate PCI	Slots 1P-5P	1
2998	Turboways 25 ATM PCI	Slots 1P-5P	4
6310	IBM ARTIC960RxD Quad Digital Trunk PCI ⁹	Slots 4P, 3P, 2P	3
6311	IBM ARTIC960RxF Quad Digital Trunk PCI ⁹	Slots 4P, 3P, 2P	3

Notes:

1. If both the POWER GXT120P or GXT130P and the POWER GXT255P or POWER GXT300P are installed, the POWER GXT255P or POWER GXT300P must be installed in slot 2 or 3.
2. The POWER GXT3000P takes the physical space of two adapters in the 7043 Model 150.
3. If two SysKonnnect SK-NET FDDI PCI Adapters are installed, they cannot be installed in slots next to each other.
4. For optimum system performance, if the 10/100 Mbps Ethernet PCI is used in 100 Mbps mode, then a maximum of one adapter is recommended per system unit.
5. For optimum system performance, when using a TURBOWAYS 155 PCI MMF ATM adapter or TURBOWAYS 155 PCI UTP ATM adapter in LAN Emulation (LANE), a maximum of one adapter per system is recommended.
6. Install only ONE of the following adapters:
 - PCI SSA Multi-Initiator/RAID EL
 - Advanced SerialRAID
 - Advanced SerialRAID Plus
7. Do not install a PCI Dual Channel Ultra2 SCSI Adapter on the same PCI bus as a POWER GXT120P or older versions of the POWER GXT130P. To determine if you have an older version of the POWER GXT130P, remove

the card and look at the barcode label on the back side. If it has a number starting with 11S94H1235, you have an older version of the card.

8. If the PCI Dual Channel Ultra2 SCSI adapter is plugged into slot 2 or 3, then the POWER GXT120P or the POWER GXT130P can only be plugged in slots 1, 4, or 5. If you plug the POWER GXT120P or the POWER GXT130P in slots 2 or 3, then you must plug the Ultra2 SCSI adapter in slot 4.
9. Do not install more than a combination of three IBM ARTIC960RxF PCI Adapters and IBM ARTIC960RxD Quad Digital Trunk PCIs in this system.
10. Do not install a PCI 3-Channel Ultra2 SCSI RAID and a IBM RTIC960RxF PCI Adapter adapter on the same bus.
11. For optimum system performance, use only two 100 Mbps ports per system.

Use Table 194 to identify adapter location combination options for these adapters in your 7043 Model 150.

Do not install more than four graphics cards in your 7043 Model 150.

Table 194. Model 150 Multiple Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2838 or 2830 or 2851 or 2852 or 2841 or 2823	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P	Slots 1P-5P Slots 1P-5P
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2845	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT550P	Slot 1P Slot 2P or 3P
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2825	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT3000P ¹	Slot 1P or 4P or 5P Slot 3P
2639 with a 2825	Ultimedia Video Capture with a GXT3000P ¹	Slot 1P, 4P, 5P Slot 3P
2639 with a 2823	Ultimedia Video Capture with a GXT2000P	Slot 1P-5P Slot 1P-5P

Feature Code	Adapter	Slot (AIX)
2639 with a 2841	Ultimedia Video Capture with a GXT300P	Slot 1P-5P Slot 1P-5P

Notes:

1. The POWER GXT3000P takes the physical space of two adapters in the 7043 Model 150.

D.7 7044 Model 170 Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 99. Model 170 System Unit Rear View with Numbered Slots

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use the table Table 195 to determine where to install an adapter in your system unit.

Use the rear-view diagram in Figure 99 on page 671 to identify slot locations described in the following tables.

Table 195. Model 170 Single Adapter Placement Guide

Feature Code	Adapter	Slot Usage	System Maximum
2825	POWER GXT3000P	Slot 1	1
2823	POWER GXT2000P	Slot 1, 2, 3, 4, 5, 6	4
2830	POWER GXT130P	Slot 3, 4, 5, 6	4
2841	POWER GXT300P	Slot 1, 2, 3, 4, 5, 6	4
2969	Gigabit Ethernet-SX PCI ²	Slot 1, 2	5
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Slot 1, 2, 3, 4, 5, 6	1
2494	PCI 3-Channel Ultra SCSI RAID ¹	Slot 3, 4, 5, 6	2
2708	ISDN Basic Rate PCI	Slot 1, 2, 3, 4, 5, 6	1
2639	Ultimedia Video Capture ¹	Slot 1, 2, 3, 4, 5, 6	1
6230	SSA Advanced SerialRAID Plus Adapter	Slot 1, 2, 3, 4, 5, 6	2
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	Slot 1, 2, 3, 4, 5, 6	2
2943	8-Port Asynchronous EIA-232/RS-422 PCI	Slot 3, 4, 5, 6	2
2944	128-Port Async Controller PCI	Slot 3, 4, 5, 6	2
2962	2-Port Multiprotocol PCI	Slot 3, 4, 5, 6	2
2963	155 Turboways ATM PCI UTP ³	Slot 1, 2, 3, 4, 5, 6	6
2988	155 Turboways ATM PCI MMF ³	Slot 1, 2, 3, 4, 5, 6	6
2742	SysKonnnect SK-NET FDDI-LP DAS PCI ⁴	Slot 1, 2, 3, 4, 5, 6	2
2741	SysKonnnect SK-NET FDDI-LP SAS PCI ⁴	Slot 1, 2, 3, 4, 5, 6	2
2743	SysKonnnect SK-NET FDDI-UP SAS PCI ⁴	Slot 1, 2, 3, 4, 5, 6	2
6205	PCI Dual Channel Ultra2 SCSI	Slot 1, 2, 3, 4, 5, 6	4
4959	4/16 Mbps token ring PCI	Slot 1, 2, 3, 4, 5, 6	4

Feature Code	Adapter	Slot Usage	System Maximum
6206	PCI Single-Ended Ultra SCSI	Slot 1, 2, 3, 4, 5, 6	4
6204	PCI Universal Differential Ultra SCSI	Slot 1, 2, 3, 4, 5, 6	4
2968	10/100 Mbps Ethernet PCI	Slot 1, 2, 3, 4, 5, 6	4
2998	Turboways 25 ATM PCI	Slot 1, 2, 3, 4, 5, 6	4

Notes:

1. Do not install PCI 3-Channel Ultra2 SCSI RAID and Ultimedia Video Capture cards on the same system.
2. For optimum system performance, install only one Gigabit Ethernet-SX PCI in this system.
3. For optimum system performance, install up to five TURBOWAYS adapters. Install up to three of these adapters if you are in MTU 1500 mode.
4. For optimum system performance, install up to five SysKonnnect adapters in this system.
5. For optimum system performance, install up to three 10/100 Mbps Ethernet PCI Adapters in this system.

Table 196. Model 170 Multiple Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)
2841 or 2823 with a 2825	GXT300P or GXT2000P with a GXT3000P	Slot 1, 2, 3, 4, 5, 6 ¹ Slot 1 ¹
2830 with a 2825	GXT130P with a GXT3000P	Slot 3, 4, 5, 6 ¹ Slot 1 ¹
2639 with a 2825	Ultimedia Video Capture with a GXT3000P	Slot 1, 2, 3, 4, 5, 6 ¹ Slot 1
2639 with a 2823	Ultimedia Video Capture with a GXT2000P	Slot 1, 2, 3, 4, 5, 6 ¹ Slot 1, 2, 3, 4, 5, 6
2639 with a 2841	Ultimedia Video Capture with a GXT300P	Slot 1, 2, 3, 4, 5, 6 ¹ Slot 1, 2, 3, 4, 5, 6

Notes:

1. Since slots 1 and 2 are capable of running at 50 MHz, the GXT2000P and GXT3000P should be placed in these slots for optimum performance. Take care when placing the GXT130P (or any other 33 MHz PCI adapter) in slots 1 and 2 since this forces the slots to run at 33 MHz.

D.8 7043 Model 260 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 100. Model 260 System Unit Rear View with Numbered Slots

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 197 on page 675 to determine where to install a single adapter in your system unit and Table 198 on page 676 to determine where to install multiple adapters in your system unit.

Use the rear view diagram in Figure 100 to identify slot locations described in the following tables.

If two different adapters can be placed in the same slot, the highest priority adapter starts at the top of the table. The list of slot numbers represent the order that the slots should be used.

Table 197. Model 260 Single Adapter Placement Guide

Feature Code	Adapter	Slot Usage	System Maximum
2838	POWER GXT120P	Slot 1, 2, 3, 4, 5	4
2830	POWER GXT130P	Slot 1, 2, 3, 4, 5	4
2851	POWER GXT250P	Slot 1, 2, 3, 4, 5	4
2852	POWER GXT255P	Slot 1, 2, 3, 4, 5	4
2841	POWER GXT300P	Slot 1, 2, 3, 4, 5	4
2823	POWER GXT2000P	Slot 1, 2, 3, 4, 5	4
2825	POWER GXT3000P	Slot 2, 4	2
6206	PCI Single-Ended Ultra SCSI	Slot 1, 2, 3, 4, 5	2
6207	PCI Differential Ultra SCSI	Slot 1, 2, 3, 4, 5	2
2494	PCI 3-Channel Ultra SCSI RAID	Slot 1, 2, 3, 4, 5	2
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	Slot 1, 2, 3, 4, 5	2
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	Slot 1, 2, 3, 4, 5	2
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	Slot 1, 2, 3, 4, 5	2
2920	Token ring PCI	Slot 1, 2, 3, 4, 5	4
2968	10/100 Mbps Ethernet PCI ¹	Slot 1, 2, 3, 4, 5	4
2969	Gigabit Ethernet-SX PCI ²	Slot 1,2	2
2708	ISDN Basic Rate PCI	Slot 1, 2, 3, 4, 5	1
2962	2-Port Multiprotocol PCI	Slot 1, 2, 3, 4, 5	2
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	Slot 1, 2, 3, 4, 5	2
2948	IBM ARTIC960Hx 4-Port T1/E1 PCI	Slot 1, 2, 3, 4, 5	2

Feature Code	Adapter	Slot Usage	System Maximum
6310	IBM ARTIC960RxD Quad Digital Trunk PCI ³	Slot 1, 2, 3, 4, 5	3
2943	8-Port Asynchronous EIA-232/RS-422 PCI	Slot 1, 2, 3, 4, 5	2
2944	128-Port Async Controller PCI	Slot 1, 2, 3, 4, 5	2
2988	155 Turboways ATM PCI MMF	Slot 1, 2, 3, 4, 5	2
2998	Turboways 25 ATM PCI	Slot 1, 2, 3, 4, 5	4
2963	155 Turboways ATM PCI UTP	Slot 1, 2, 3, 4, 5	2
6215	PCI SSA Multi-Initiator/RAID EL	Slot 1, 2, 3, 4, 5	1
6225	Advanced SerialRAID	Slot 1, 2, 3, 4, 5	2
2639	Ultimedia Video Capture	Slot 1, 2, 3, 4, 5	1

Notes:

1. For optimum system performance, if the 10/100 Mbps Ethernet PCI is used in 100 Mbps mode, a maximum of three adapters is recommended per system unit.
2. For optimum system performance, a maximum of 1 Gigabit Ethernet PCI is recommended per system unit.
3. See Appendix D.19, "Digital Trunk PCI Adapter Placement Considerations" on page 709.

Use Table 198 to identify adapter location combination options for these adapters in your 7043 Model 260.

Do not install more than four graphics cards in your 7043 Model 260.

Table 198. Model 260 Multiple Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2838 or 2830 or 2851 or 2852 or 2841 or 2823	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P	Slot 1-5 ¹ Slot 1-5 ¹

Feature Code	Adapter	Slot (AIX)
2838 or 2830 or 2851 or 2852 or 2841 or 2823 with a 2825	GXT120P or GXT130P or GXT250P or GXT255P or GXT300P or GXT2000P with a GXT3000P	Slot 1, 3, 4, 5 ¹ Slot 2 ¹
2639 with a 2825	Ultimedia Video Capture with a GXT3000P	Slot 1, 3, 4, 5 ¹ Slot 2
2639 with a 2823	Ultimedia Video Capture with a GXT2000P	Slot 1-5 ¹ Slot 1-5
2639 with a 2841	Ultimedia Video Capture with a GXT300P	Slot 1-5 ¹ Slot 1-5

Notes:

1. Since slots 1 and 2 are capable of running at 50 mhz, the GXT250P, GXT255P, GXT2000P, GXT 3000P, and POWER GXT300P should be placed in these slots for optimum performance. Take care when placing the GXT120P or GXT130P (or any other 33 MHz PCI adapter) in slots 1 and 2 since this forces the slots to run at 33 MHz as a POWER GXT120P or a POWER GXT130P.

D.9 7043/7044 Model 270 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 101. Model 270 System Unit Rear View with Numbered Slots

Use Table 199 to identify specific slot location options for the following adapters in your Machine Types 7043, 7044 Model 270 system and Table 200 on page 680 to determine where to install multiple adapters in your system unit.

If two different adapters can be placed in the same slot, the highest priority adapter starts at the top of the table. The list of slot numbers represent the order that the slots should be used.

Use the rear view diagram in Figure 101 to identify slot locations described in the following tables.

Table 199. Model 270 Single Adapter Placement Guide

Feature Code	Adapter	Slot Usage	System Maximum
2825	POWER GXT3000P	Slot 2, 4	2
2851	POWER GXT250P	Slot 1, 2, 3, 4, 5	4
2852	POWER GXT255P	Slot 1, 2, 3, 4, 5	4
2823	POWER GXT2000P	Slot 1, 2, 3, 4, 5	4

Feature Code	Adapter	Slot Usage	System Maximum
2838	POWER GXT120P	Slot 1, 2, 3, 4, 5	4
2830	POWER GXT130P	Slot 1, 2, 3, 4, 5	4
2841	POWER GXT300P	Slot 1, 2, 3, 4, 5	4
2969	Gigabit Ethernet-SX PCI ²	Slot 1, 2	5
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Slot 1, 2, 3, 4, 5	1
2494	PCI 3-Channel Ultra SCSI RAID ¹	Slot 3, 4, 5	2
2708	ISDN Basic Rate PCI	Slot 3, 4, 5	1
2639	Ultimedia Video Capture ¹	Slot 3, 4, 5	5
6215	PCI SSA Multi-Initiator/RAID EL	Slot 3, 4, 5	1
6230	SSA Advanced SerialRAID Plus Adapter	Slot 1, 2, 3, 4, 5	2
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	Slot 1, 2, 3, 4, 5	2
2943	8-Port Asynchronous EIA-232/RS-422 PCI	Slot 1, 2, 3, 4, 5	2
2944	128-Port Async Controller PCI	Slot 1, 2, 3, 4, 5	2
2962	2-Port Multiprotocol PCI	Slot 3, 4, 5	2
2963	155 Turboways ATM PCI UTP	Slot 1, 2, 3, 4, 5	5
2988	155 Turboways ATM PCI MMF	Slot 1, 2, 3, 4, 5	5
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	Slot 1, 2, 3, 4, 5	5
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	Slot 1, 2, 3, 4, 5	5
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	Slot 1, 2, 3, 4, 5	2
6205	PCI Dual Channel Ultra2 SCSI	Slot 1, 2, 3, 4, 5	2
4959	4/16 Mbps token ring PCI	Slot 1, 2, 3, 4, 5	4
6206	PCI Single-Ended Ultra SCSI	Slot 3, 4, 5	3
6207	PCI Differential Ultra SCSI	Slot 1, 2, 3, 4, 5	4
2968	10/100 Mbps Ethernet PCI ³	Slot 1, 2, 3, 4, 5	5
2998	Turboways 25 ATM PCI	Slot 1, 2, 3, 4, 5	4
6204	PCI Universal Differential Ultra SCSI	Slot 1, 2, 3, 4, 5	4

Notes:

1. Do not install PCI 3-Channel Ultra2 SCSI RAID and Ultimedia Video Capture cards on the same system.
2. For optimum system performance, install up to two Gigabit Ethernet-SX PCI adapters per system. If you install more Gigabit Ethernet-SX PCI adapters, put them in slots 3, 4, and 5, respectively.
3. Machine Types 7043, 7044 Model 270 support up to 6 10/100 Mbps Ethernet PCI Adapters per system if you use the Integrated Ethernet port.
4. Do not install more than four graphics adapters in your Machine Types 7043, 7044 Model 270.

Table 200. Model 270 Multiple Adapter Placement Guide

Feature Code	Adapter	Slot (AIX)
2830 or 2841 or 2823 with a 2825	GXT130P or GXT300P or GXT2000P with a GXT3000P	Slot 1, 2, 3, 4, 5, 6 ¹ Slot 1, 2, 3, 4, 5 ¹
2825 with a 2825	GXT3000P with a GXT3000P	Slot 2 ¹ Slot 4 ¹
2639 with a 2825	Ultimedia Video Capture with a GXT3000P	Slot 1, 3, 4, 5 ¹ Slot 2
2639 with a 2823	Ultimedia Video Capture with a GXT2000P	Slot 1, 3, 4, 5 ¹ Slot 2
2639 with a 2841	Ultimedia Video Capture with a GXT300P	Slot 1, 2, 3, 4, 5 ¹ Slot 1, 2, 3, 4, 5

Notes:

1. Since slots 1 and 2 are capable of running at 50 MHz, the GXT2000P and GXT3000P should be placed in these slots for optimum performance. Take care when placing the GXT130P (or any other 33 MHz PCI adapter) in slots 1 and 2 since this forces the slots to run at 33 MHz.

D.10 7046 Model B50 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.




Figure 102. Model B50 System Unit Rear View with Numbered Slots

Table 201 summarizes the adapters (communication, storage, and graphics) available on the RS/6000 7046 Model B50 system and describes the adapter placement guidelines.

Table 201. Communication, Graphics, and Storage Adapters Summary

Adapter	Max per System	Slot Allocation	Feature Code	Size
4/16 High-Speed token ring PCI	2	1-2	# 4959	short
10/100 Mbps Ethernet PCI	2	1-2	# 2968	short
ISDN Basic Rate PCI	1	1-2	# 2708	short
SysKonnect SK-NET FDDI-LP SAS PCI	1	1-2	# 2741	short
SysKonnect SK-NET FDDI-LP DAS PCI	1	1-2	# 2742	short
SysKonnect SK-NET FDDI-UP SAS PCI	1	1-2	# 2743	short
TURBOWAYS 155 UTP ATM PCI	2	1-2	# 2963	short
TURBOWAYS 155 MMF ATM PCI	2	1-2	# 2988	short
8-Port Asynchronous EIA-232/RS-422 PCI	2	1-2	# 2943	short
128-Port Asynchronous Controller PCI	2	1-2	# 2944	short
2-Port Multiprotocol PCI	2	1-2	# 2962	short
ARTIC960Hx 4-Port Selectable PCI	1	2	# 2947	long
POWER GXT130P	1	1	# 2830	short
PCI Single-Ended Ultra SCSI	2	1-2	# 6206	short
PCI Differential Ultra SCSI	2	1-2	# 6207	short
PCI Dual-Channel Ultra2 SCSI Adapter	1	1-2	# 6205	short
PCI 3-Channel Ultra2 SCSI RAID	1	2	# 2494	long

Adapter	Max per System	Slot Allocation	Feature Code	Size
PCI Advanced SerialRAID	1	2	# 6225	long

D.11 7025 Model F50 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.

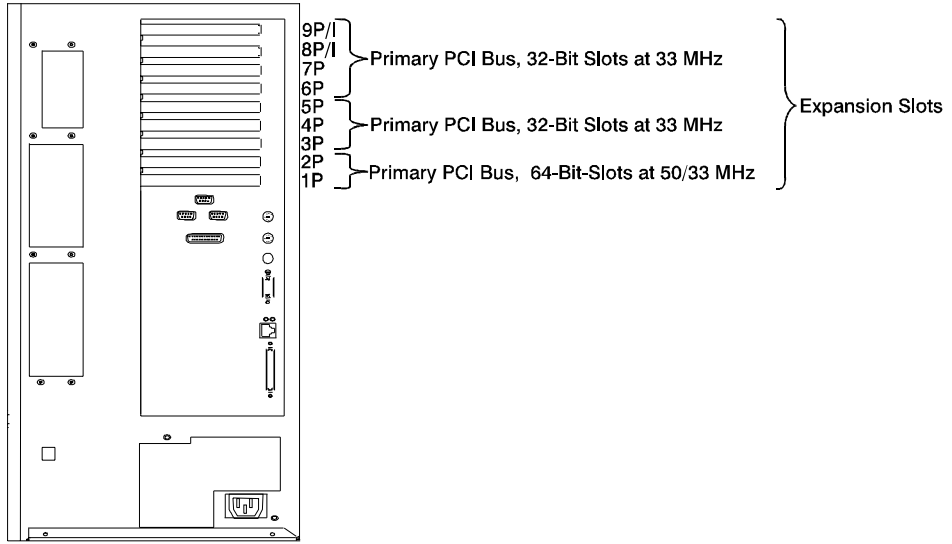


Figure 103. Model F50 System Unit Rear View with Numbered Slots

Note

The Machine Type 7025 Model F50 has three PCI buses, with Slots 1 - 2 on one bus, Slots 3 - 5 on a bus, and Slots 6 - 9 on a bus.

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 202 on page 683 to determine where to install these adapters in your 7025 Model F50.

Use the rear view diagram in Figure 103 on page 682 to identify slot locations described in the following table.

Table 202. Model F50 Adapter Placement Guide

Feature Code	Adapter	Usage
2969	Gigabit Ethernet PCI ⁶	-Max. of 1 in slots 3P, 4P, or 5P -Max. of 1 in slots 6P, 7P, 8P/I, or 9P/I
6310	IBM ARTIC960RxD Quad Digital Trunk PCI ⁸	Any slot (Max. 4 per system)
2839	POWER GXT110P ¹	Slot 3P, 4P, or 5P
2838	POWER GXT120P ¹	Any slot
2853	POWER GXT800P ¹	Slot 3P ⁹ (Max. 1 per system)
6215	PCI SSA Multi-Initiator/RAID EL ⁵	Max. 4 per system, 2 per bus (see note...)
6225	Advanced SerialRAID ⁵	Max. 4 per system, 2 per bus (see note...)
6309 2751	Digital Trunk Quad PCI ⁸ S/390 ESCON Channel PCI	If less than 2 GB of system memory is installed, a maximum of 3 adapters combined total per system: -Max. of 3 in slots 1P, 2P, 3P, 4P, or 5P. -Max. of 2 in slots 6P, 7P, 8P/I, or 9P/I. If more than 2 GB of system memory is installed, a maximum of 3 adapters combined total per system: -Max. of 1 in slots 1P or 2P -Max. of 2 in slots 3P, 4P, or 5P -Max. of 1 in slots 6P, 7P, 8P/I, or 9P/I.

Notes:

1. The 7025 Model F50 supports combinations of the POWER GXT110P, POWER GXT120P, and the POWER GXT800P Adapters up to a maximum of two adapters total per system.
2. The 7025 Model F50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of four per system. For best performance, the combination should not exceed two per PCI bus:

- 155 Turboways ATM PCI MMF
 - 155 Turboways ATM PCI UTP
 - 10/100 Mbps PCI Fast EtherLink XL
 - 10/100 Mbps Ethernet PCI
 - SCSI-2 Fast/Wide PCI RAID
 - SysKonnnect SK-NET FDDI-LP DAS PCI
 - SysKonnnect SK-NET FDDI-LP SAS PCI
 - SysKonnnect SK-NET FDDI-UP SAS PCI
3. The 7025 Model F50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of seven:
 - IBM ARTIC960Hx 4-Port Selectable PCI
 - IBM ARTIC960Hx 4-Port T1/E1 PCI
 - IBM ARTIC960Hx DSP Resource PCI
 - S/390 ESCON Channel PCI
 - Digital Trunk Quad PCI
 4. The 7025 Model F50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of three:
 - Digital Trunk Quad PCI
 - S/390 ESCON Channel PCI
 5. For optimum system performance, the combination of the following adapters should not exceed a maximum of one adapter per bus:
 - PCI SSA Multi-Initiator/RAID EL
 - Advanced SerialRAID
 - PCI 3-Channel Ultra SCSI RAID
 - Gigabit Ethernet PCI
 6. For optimum system performance, a maximum of one adapter per system is recommended.
 7. The 7025 Model F50 supports combinations of the following adapters, but that COMBINATION cannot exceed a total of four.
 - PCI SSA Multi-Initiator/RAID EL
 - Advanced SerialRAID
 8. See Appendix D.19, "Digital Trunk PCI Adapter Placement Considerations" on page 709.

9. Slots 1P, 2P, and 4P cannot be used when a GXT800P is installed in slot 3P.

D.12 7025 Model F80 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing. Many of the following notes refer to optimizing system performance.

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use the rear-view diagram shown in Figure 104 to identify slot locations described in the Table 203.




Figure 104. Model F80 System Unit Rear View with Numbered Slots

The adapters with the highest slot placement priority are listed at the top of the table. The slot numbers in the Slot Usage column represent slot location priorities. Use the first numbered slot first. If an adapter has already filled the first slot in the list, go to the next number in the list.

Table 203. Model F80 Adapter Placement Guide

Feature Code	Adapter	Usage	F80 Maximum	Hot-Pluggable?
2830	POWER GXT130P	6, 11, 7, 12	1	No
2969	Gigabit Ethernet-SX PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Yes

Feature Code	Adapter	Usage	F80 Maximum	Hot-Pluggable?
6204	PCI Differential Ultra SCSI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Yes
6208	PCI SCSI-2 Single Ended Fast/Wide	6, 11, 7, 12	4	Yes
6205	PCI Dual Channel Ultra2 SCSI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Yes
2494	PCI 3-Channel Ultra2 SCSI RAID	6, 11, 7, 12	4	Yes
2742	SysKonnnect SK-NET FDDI-LP DAS PCI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Yes
2741	SysKonnnect SK-NET FDDI-LP SAS PCI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Yes
2743	SysKonnnect SK-NET FDDI-UP SAS PCI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Yes
4959	High-Speed token ring PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	9	Yes
2968	10/100 Mbps Ethernet PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Yes
2985	Ethernet 10base2 PCI	6, 11, 7, 12	4	Yes
2708	ISDN Basic Rate PCI	6, 11, 7, 12	1	Yes
2962	2-Port Multiprotocol PCI	6, 11, 7, 12	4	Yes
2947	IBM Artic960Hx 4-Port Multiprotocol PCI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	No
2943	8-Port Asynchronous EIA-232E/RS-422A PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Yes
2944	128-Port Async Controller PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Yes
2963	TURBOWAYS 155 PCI UTP ATM	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Yes
2988	TURBOWAYS 155 PCI MMF ATM	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Yes ¹
6225	Advanced SerialRAID	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (2/bus)	Yes
6230	Advanced SerialRAID Plus	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Yes

Feature Code	Adapter	Usage	F80 Maximum	Hot-Pluggable?
2751	S/390 ESCON Channel PCI	6, 11, 7, 12	3	No
6227	Gigabit Fiber Channel for PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Yes
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	See note 1 below this table
6311	IBM ARTIC960RxF PCI Adapter	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	No

Notes:

1. This adapter is hot-pluggable if it is not cabled to another card internally.

D.13 7026 Model H50 Adapter Placement Guide

This system is designed for service representatives to install adapters. Use this guide to determine if there are specific slot requirements for adapters that your service representative may be installing.




Figure 105. Model H50 System Unit Rear View with Numbered Slots

Note

The Machine Type 7026 Model H50 has three PCI buses, with Slots 1 - 2 on one bus, Slots 3 - 5 on a bus, and Slots 6 - 9 on a bus.

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 204 to determine where to install these adapters in your 7026 Model H50.

Use the rear view diagram in Figure 105 on page 687 to identify slot locations described in Table 204.

Table 204. Model H50 Adapter Placement Guide

Feature Code	Adapter	Usage
2969	Gigabit Ethernet PCI ⁶	-Max. of 1 in slots 3P, 4P, or 5P -Max. of 1 in slots 6P, 7P, 8P/I, or 9P/I
6310	IBM ARTIC960RxD Quad Digital Trunk PCI ⁷	Any slot (Max. 4 per system)
2838	POWER GXT120P	Any slot (Max. 2 per system)
6215	PCI SSA Multi-Initiator/RAID EL ⁵	Slots 1P, 4P, 7P, 9P/I
6225	Advanced SerialRAID ⁵	Slots 1P, 4P, 7P, 9P/I
6309 2751	Digital Trunk Quad PCI ⁷ S/390 ESCON Channel PCI	If less than 2 GB of system memory is installed, a maximum of 3 adapters combined total per system: -Max. of 3 in slots 1P, 2P, 3P, 4P, or 5P. -Max. of 2 in slots 6P, 7P, 8P/I, or 9P/I. If more than 2 GB of system memory is installed, a maximum of 3 adapters combined total per system: -Max. of 1 in slots 1P or 2P -Max. of 2 in slots 3P, 4P, or 5P -Max. of 1 in slots 6P, 7P, 8P/I, or 9P/I.

Notes:

1. A PCI SSA Multi-Initiator/RAID EL Adapter MAY NOT be plugged into a slot adjacent to a slot occupied by a full-length adapter.

2. The 7026 Model H50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of four per system. For best performance, the COMBINATION should not exceed two of these adapters per PCI bus:
 - 155 Turboways ATM PCI MMF
 - 155 Turboways ATM PCI UTP
 - 10/100 Mbps PCI Fast EtherLink XL
 - 10/100 Mbps Ethernet PCI
 - SCSI-2 Fast/Wide PCI RAID
 - SysKonnnect SK-NET FDDI-LP DAS PCI
 - SysKonnnect SK-NET FDDI-LP SAS PCI
 - SysKonnnect SK-NET FDDI-UP SAS PCI
3. The 7026 Model H50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of seven:
 - IBM ARTIC960Hx 4-Port Selectable PCI
 - IBM ARTIC960Hx 4-Port T1/E1 PCI
 - IBM ARTIC960Hx DSP Resource PCI
 - S/390 ESCON Channel PCI
 - Digital Trunk Quad PCI
4. The 7026 Model H50 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of three:
 - Digital Trunk Quad PCI
 - S/390 ESCON Channel PCI
5. For optimum system performance, the combination of the following adapters should not exceed a maximum of one adapter per bus.
 - PCI SSA Multi-Initiator/RAID EL
 - Advanced SerialRAID
 - PCI 3-Channel Ultra SCSI RAID
 - Gigabit Ethernet PCI
6. For optimum system performance, a maximum of one adapter per system is recommended.
7. See Appendix D.19, “Digital Trunk PCI Adapter Placement Considerations” on page 709.

D.14 7026 Model H70 Adapter Placement Guide

This system is designed for service representatives to install adapters. Use this guide to determine if there are specific slot requirements for adapters that your service representative may be installing.




Figure 106. Model H70 System Unit Rear View with Numbered Slots

Some adapters must be placed in specific system unit slots to function correctly at highest performance. Use Table 205 to determine where to install an adapter in your system unit.

Use the rear view diagram in Figure 106 to identify slot locations described in the following tables.

Table 205. Model H70 Adapter Placement Guide

Feature Code	Adapter	Usage	System Maximum
2948	IBM ARTIC960HX 4-Port T1/E1 PCI ²	Any slot	7
2947	IBM ARTIC960Hx 4-Port Selectable PCI ²	Any slot	7

Feature Code	Adapter	Usage	System Maximum
2949	IBM ARTIC960Hx DSP Resource PCI ²	Any slot	7
6215	PCI SSA Multi-Initiator/RAID EL ⁵	Any slot	4
2920	Token ring PCI ^{3,5}	Any slot	8
6310	IBM ARTIC960RxD Quad Digital Trunk PCI ^{2,7}	Any slot	4
2969	Gigabit Ethernet PCI ^{3,6}	Slots 1,2 5, 7	4
6225	Advanced SerialRAID ³	Any slot	4
2751	S/390 ESCON Channel PCI ^{2,4}	Slots 3, 4, 5, 6, 7, 8	3
2494	PCI 3-Channel Ultra SCSI RAID ³	Any slot	4
2943	8-Port Asynchronous Adapter EIA-232E/RS-422A PCI	Any slot	8
2944	128-Port Async Controller PCI	Any slot	8
6206	PCI Ultra SCSI Single-Ended	Any slot	8
6207	PCI Differential Ultra SCSI	Any slot	8
6208	PCI SCSI-2 Fast/Wide Single Ended	Any slot	8
6209	PCI SCSI-2 Fast/Wide Differential	Any slot	8
2830	POWER GXT130P	Any slot	2
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	Any slot	8
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	Any slot	8
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	Any slot	8
2963	155 TURBOWAYS ATM PCI UTP ¹	Any slot	8
2968	10/100 Mbps Ethernet PCI ¹	Any slot	8
2979	PCI Auto LANStreamer token ring	Any slot	8
2985	Ethernet 10base2 PCI	Any slot	8

Feature Code	Adapter	Usage	System Maximum
2987	Ethernet 10base5 PCI	Any slot	8
2708	ISDN Basic Rate PCI	Any slot	1
2988	155 TURBOWAYS ATM PCI MMF	Any slot	8
2962	2-Port Multiprotocol PCI	Any slot	8
6309	Digital Trunk Quad PCI ^{4,7}	Any slot	3

Notes:

1. For optimum performance, the 7026 Model H70 supports any combination of the following adapters, but that COMBINATION should not exceed a total of six per system:
 - 155 Turboways ATM PCI MMF when used in LAN emulation (LANE) mode.
 - 155 Turboways ATM PCI UTP when used in LAN emulation (LANE) mode.
 - 10/100 Mbps Ethernet PCI when used in 100 Mbps mode. If the integrated 10/100 Mbps controller is also being used at 100 Mbps, the number of 10/100 Mbps Ethernet PCI should be limited to five.
2. The 7026 Model H70 supports any combination of the following adapters, but that COMBINATION cannot exceed a total of seven:
 - IBM ARTIC960Hx 4-Port T1/E1 PCI
 - IBM ARTIC960Hx 4-Port Selectable PCI
 - IBM ARTIC960Hx DSP Resource PCI
 - S/390 ESCON Channel PCI
 - Digital Trunk Quad PCI
3. For optimum performance, the combination of the following adapters should not exceed one per PCI bus:
 - Gigabit Ethernet PCI
 - Advanced SerialRAID
 - PCI 3-Channel Ultra SCSI RAID
 - PCI SSA Multi-Initiator/RAID EL
4. The 7026 Model H70 supports any combination of the following adapters, but only one of these adapters can be installed in slots 5 and 6:

- S/390 ESCON Channel PCI
 - Digital Trunk Quad PCI
5. A PCI SSA Multi-Initiator/RAID EL adapter MAY NOT be plugged into a slot adjacent to a slot occupied by a full-length adapter.
 6. For optimum performance, a maximum of two Gigabit Ethernet PCI Adapters per system is recommended.
 7. See Appendix D.19, “Digital Trunk PCI Adapter Placement Considerations” on page 709.

D.15 7026 Model H80 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.

Many of the following notes refer to optimizing system performance.

Use the rear-view diagram shown in Figure 107 to identify slot locations described in the Table 206 on page 694.




Figure 107. Model H80 System Unit Rear View with Numbered Slots

The adapters with the highest slot placement priority are listed at the top of the table. The slot numbers in the Slot Usage column represent slot location

priorities. Use the first numbered slot first. If an adapter has already filled the first slot in the list, go to the next number in the list.

Table 206. Model H80 Adapter Placement Guide

Feature Code	Adapter	Usage	Model H80 Maximum	Hot-Pluggable?
2830	POWER GXT130P	Primary 6, 11, 7, 12	1	Yes (see note 2 below this table)
2838	POWER GXT120P	Primary 6, 11, 7, 12	1	Yes (see note 2 below this table)
6204	PCI Universal Differential Ultra SCSI Adapter	Primary drawer 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	10	Yes (see note 2 below this table)
6208	PCI SCSI-2 Single Ended Fast/Wide	Primary drawer 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	4	Yes (see note 2 below this table)
6207	PCI Differential Ultra SCSI Adapter	Primary drawer 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	4	Yes (see note 2 below this table)
6205	PCI Dual Channel Ultra2 SCSI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	10	Yes (see note 2 below this table)
2494	PCI 3-Channel Ultra2 SCSI RAID	Primary 6, 11, 7, 12	4	Yes
2742	SysKonnnect SK-NET FDDI-LP DAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2741	SysKonnnect SK-NET FDDI-LP SAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2743	SysKonnnect SK-NET FDDI-UP SAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes

Feature Code	Adapter	Usage	Model H80 Maximum	Hot-Pluggable?
4959	High-Speed token ring PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2920	PCI token ring	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	8	Yes
2979	PCI Auto LANstreamer token ring	Primary 6, 11, 7, 12	4	Yes
2968	10/100 Mbps Ethernet PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2969	Gigabit Ethernet-SX PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	10	Yes
2708	ISDN Basic Rate PCI	Primary 6, 11, 7, 12	1	Yes
2962	2-Port Multiprotocol PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 see note 3 below this table	5	Yes
2947	IBM Artic960Hx 4-Port Multiprotocol PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	No
2943	8-Port Asynchronous EIA-232E/RS-422A PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2944	128-Port Async Controller PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2963	TURBOWAYS 155 PCI UTP ATM	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes
2988	TURBOWAYS 155 PCI MMF ATM	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	Yes ¹
6225	Advanced SerialRAID	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	4 (2/bus)	Yes (see note 2 below this table)
6230	Advanced SerialRAID Plus	Primary 6, 11, 7, 12	8	Yes (see note 2 below this table)
6215	PCI SSA Multi-Initiator/RAID EL	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	4	No

Feature Code	Adapter	Usage	Model H80 Maximum	Hot-Pluggable?
2751	S/390 ESCON Channel PCI	Primary 6, 11, 7, 12	2	No
6227	Gigabit Fiber Channel for PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	8	Yes
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Primary 9, 10, 11, 12, 13, 14	4	see note 1 on below this table

Notes:

1. Do not hot-plug this adapter when it is internally cabled to another adapter.
2. Do not hot-plug any PCI adapter supporting the system's boot device or system console.
3. Adapter migration may require the use of a 5V slot.

D.16 7026 Model M80 Adapter Placement Guide

This system is designed for customers to install adapters. Use this guide to determine if there are specific slot requirements for adapters that you may be installing.

Many of the following notes refer to optimizing system performance.

Use the rear-view diagram shown in Figure 108 on page 697 to identify slot locations described in the Table 207.




Figure 108. Model M80 System Unit Rear View with Numbered Slots

The adapters with the highest slot placement priority are listed at the top of the table. The slot numbers in the Slot Usage column represent slot location priorities. Use the first numbered slot first. If an adapter has already filled the first slot in the list, go to the next number in the list.

Table 207. Model M80 Adapter Placement Guide

Feature Code	Adapter	Usage	Model M80 Maximum	Hot-Pluggable?
2830	POWER GXT130P	Primary drawer slot 6, 11, 7, 12 Secondary drawer slot 6, 11, 7, 12	1 per system	Yes (see note 2 below this table)
6204	PCI Universal Differential Ultra SCSI Adapter	Primary drawer slot 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary drawer slot 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	40 per system	Yes (see note 2 below this table)
6205	PCI Dual Channel Ultra2 SCSI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	20 per system	Yes (see note 2 below this table)
2494	PCI 3-Channel Ultra2 SCSI RAID	Primary 6, 11, 7, 12 Secondary 6, 11, 7, 12	8 per drawer	Yes
2742	SysKonnnect SK-NET FDDI-LP DAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28 per drawer	Yes

Feature Code	Adapter	Usage	Model M80 Maximum	Hot-Pluggable?
2741	SysKonnnect SK-NET FDDI-LP SAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28	Yes
2743	SysKonnnect SK-NET FDDI-UP SAS PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28	Yes
4959	High-Speed token ring PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28	Yes
2968	10/100 Mbps Ethernet PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28 per system	Yes
2969	Gigabit Ethernet-SX PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14	20 per system	Yes
2708	ISDN Basic Rate PCI	Primary 6, 11, 7, 12 Secondary 6, 11, 7, 12	1 per system	Yes
2962	2-Port Multiprotocol PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	20 per system	Yes
2947	IBM Artic960Hx 4-Port Multiprotocol PCI Adapter	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	14	No
2943	8-Port Asynchronous EIA-232E/RS-422A PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	32 per system	Yes
2944	128-Port Async Controller PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	32 per system ⁶	Yes

Feature Code	Adapter	Usage	Model M80 Maximum	Hot-Pluggable?
2963	TURBOWAYS 155 PCI UTP ATM	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28 per system ⁷	Yes
2988	TURBOWAYS 155 PCI MMF ATM	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	28 per system ⁸	Yes
6225	Advanced SerialRAID	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	4 per system	Yes (see note 2 below this table)
6230	Advanced SerialRAID Plus	Primary 6, 11, 7, 12 Secondary 6, 11, 7, 12	32 per system	Yes (see note 2 below this table)
2751	S/390 ESCON Channel PCI	Primary 6, 11, 7, 12 Secondary 6, 11, 7, 12 (see note 3 below this table)	4 per system	No
6227	Gigabit Fiber Channel for PCI	Primary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12 Secondary 1, 5, 10, 2, 8, 4, 13, 3, 9, 14, 6, 11, 7, 12	32 per system	Yes
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Primary 9, 10, 11, 12, 13, 14 Secondary 9, 10, 11, 12, 13, 14	4 per system	see note 1 on below this table

Notes:

1. Do not hot-plug this adapter when it is internally cabled to another adapter.
2. Do not hot-plug any PCI adapter supporting the system's boot device or system console.
3. If you have more than two drawers, do not install the S/390 ESCON Channel PCI in drawer three or higher.

D.17 7017 Model S70, S70 Advanced, and S80 Adapter Placement Guide

These systems are designed for service representatives to install adapters. Use this guide to determine if there are specific slot requirements for adapters that your service representative may be installing.




Figure 109. I/O Drawer Rear View with 14 Slots and 4 PCI Buses

Note

Each drawer has four PCI buses per Drawer: Slots 1-4 (PCI Bus 1), 5-8 (PCI Bus 0), 9-10 (PCI Bus 3), 11-14 (PCI Bus 2). Slots 1, 5, 9, 10, and 14 are 64-bit slots. Remaining slots are 32-bit. The 32-bit adapters also function in the 64-bit slots. All slots are 33 MHz.

D.17.1 Adapter Placement

Some adapters must be placed in specific I/O drawer slots to function correctly and to achieve highest performance.

Many of the following notes refer to optimizing system performance. Read "System Performance" on page 1 for more performance-related information.

The following table shows the slot plugging guidelines for Models S70, S7A, and S80 systems. The information in these tables has been verified in system testing performed using maximum configurations.

If two different adapters can be placed in the same slot, the highest priority adapter starts at the top of the table. The list of slot numbers represent the order that the slots should be used in a drawer.

For complete system placement, the first slot for a group of adapters is selected in the first drawer, and then the first slot is selected in the next drawer. After the first slot selection has been used for each drawer, the second slot in the list is used. This also rotates through the available drawers. If a card has already been placed in a slot, the slot is not available for future adapter placement.

The primary drawer referenced in Table 208 is the I/O drawer that has a service processor card installed in slot 8. Secondary drawers are additional I/O drawers that may be added to a Models S70, S7A, and S80 configuration. Secondary drawers do not have a service processor installed.

Table 208. Models S70, S70 Advanced Adapter, and S80 Placement Guide

Adapter	Slot Usage	S70 Max	S7A Max	S80 Max
Service Processor	Primary drawer slot 8	1 per system	1 per system	1 per system
SP System Attachment PCI (see *9076 RS/6000 SP Systems)	Primary drawer slot 10, slots 9 and 11 must be empty	1 per system	1 per system	1 per system
S/390 ESCON Channel PCI	Primary drawer slots 10,14 Secondary drawer slots 10,14 S/390 ESCON Channel PCI adapters can only be installed in the primary I/O drawer and the secondary I/O drawer that are cabled in the same RIO loop (see note 11)	4 per system	4 per system	4 per system
PCI SCSI-2 F/W Single-Ended Factory installed to support internal media drives	S70 Only Primary drawer - slot 8 (as required) Systems manufactured before October 23, 1998 may have this adapter installed in slot 2	1 per drawer	N/A	N/A
PCI SCSI-2 F/W Single-Ended Factory installed to support internal SCSI drives	S70 Only Primary drawer - slot 13 Secondary drawer - slot 13 (as required) Systems manufactured before October 23, 1998 may have this adapter installed in slot 9.	1 per drawer	N/A	N/A
PCI Ultra SCSI Single-Ended Factory installed to support internal media drives.	S7A and S80 Only Primary drawer - slot 7 Secondary drawer - slot 8 (As required)	N/A	1 per drawer	1 per drawer

Adapter	Slot Usage	S70 Max	S7A Max	S80 Max
PCI Ultra SCSI Single-Ended Factory installed to support internal SCSI drives	S7A and S80 Only Primary drawer - slot 13, 6 Secondary drawer - slot 13, 6 (As required)	N/A	2 per drawer	N/A
Gigabit Fibre Channel for PCI (see note 5)	Primary drawer - slots 10, 14, 1 Secondary drawer - slots 1, 5 10, 14	4 per system 4 per I/O drawer	15 per system 4 per I/O drawer	15 per system 4 per I/O drawer
Gigabit Ethernet-SX PCI	Primary drawer - slots 10, 1, 9, 3 Secondary drawer - slots 1, 5, 10, 9, 3	8 per system (see Performan ce Limits)	8 per system (see Performan ce Limits)	19 per system (see Performan ce Limits)
PCI SSA Multi-Initiator/RAID EL (see note 1 and 2)	Primary drawer - slots 3, 14, 10, 12, 1 Secondary drawer - slots 1, 5, 10, 14, 3, 7, 12	26 per system	26 per system	26 per system
Advanced SerialRaid SSA Plus (see note 2)	Primary drawer - slots 3, 14, 10, 12, 1 Secondary drawer - slots 1, 5, 10, 14, 3, 7, 12 (see note 12)	26 per system	26 per system	26 per system
Advanced SerialRAID (see note 2)	Primary drawer - slots 3, 14, 10, 12, 1 Secondary drawer - slots 1, 5, 10, 14, 3, 7, 12	26 per system	26 per system	26 per system
PCI Dual Channel Ultra2 SCSI	Primary drawer - slots 3, 14, 10, 12, 1 Secondary drawer - slots 1, 5, 10, 14	N/A	26 per system (see note 10)	26 per system (see note 10)
POWER GXT120P (see note 3)	Primary drawer - slots 1, 4 (see note 12)	1 per system	1 per system	1 per system
POWER GXT130P (see note 3)	Primary drawer - slots 1, 4 (see note 12)	1 per system	1 per system	1 per system
155 TURBOWAYS ATM PCI UTP 155 TURBOWAYS ATM PCI MMF (see note 8)	Primary drawer - slots 2, 4, 9, 11, 13, 3, 10, 12, 6, 1, 5 Secondary drawer - slots 2, 4, 6, 8, 9, 11, 13, 3, 7, 10, 12, 14, 1, 5 (see note 11)	16 per system (see note 6)	16 per system (see note 6)	16 per system (see note 6)

Adapter	Slot Usage	S70 Max	S7A Max	S80 Max
10/100 Mbps Ethernet PCI (see note 8)	Primary drawer - slots 2, 4, 9, 11, 13, 3, 10, 12, 6, 1, 5 Secondary drawer - slots 2, 4, 6, 8, 9, 11, 13, 3, 7, 10, 12, 14, 1, 5 (see note 11)	26 per system (see note 6)	26 per system (see note 6)	26 per system (see note 6)
High-Speed PCI token ring	Primary drawer - slots 2, 4, 9, 11, 13, 3, 10, 12, 6, 1, 5 Secondary drawer - slots 2, 4, 6, 8, 9, 11, 13, 3, 7, 10, 12, 14, 1, 5	26 per system	26 per system	26 per system
SysKonnnect SK-NET FDDI-UP SAS PCI SysKonnnect SK-NET FDDI-LP DAS PCI SysKonnnect SK-NET FDDI-LP SAS PCI	Primary drawer - slots 2, 4, 9, 11, 13, 3, 10, 12, 6, 1, 5 Secondary drawer - slots 2, 4, 6, 8, 9, 11, 13, 3, 7, 10, 12, 14, 1, 5	12 per system (see note 6)	26 per system (see note 6)	26 per system (see note 6)
IBM ARTIC960Hx 4-Port Selectable PCI	Primary drawer - slots 1, 3, 5, 10, 12, 14 Secondary drawer - slots 1, 3, 5, 7, 10, 12, 14	7 per system	14 per system	14 per system
IBM ARTIC960Hx 4-Port T1/E1 PCI	Primary drawer - slots 1, 3, 5, 10, 12, 14 Secondary drawer - slots 1, 3, 5, 7, 10, 12, 14	14 per system	14 per system	14 per system
IBM ARTIC960RxD Quad Digital Trunk PCI	Primary drawer - slots 1, 3, 5, 10, 12, 14 Secondary drawer - slots 1, 3, 5, 7, 10, 12, 14 (see Chapter "Digital Trunk PCI Adapter Placement Considerations and note 13)	4 per system	4 per system	4 per system
IBM ARTIC960RxF PCI Adapter	Any slot (see Chapter "Digital Trunk PCI Adapter Placement Considerations and note 13)	N/A	3	N/A

Adapter	Slot Usage	S70 Max	S7A Max	S80 Max
8-Port Asynchronous PCI	Primary drawer - slots 1, 3, 4, 11, 12, 13, 14, 10, 5, 9 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 14, 2, 6, 9, 13, 4, 7	8 per system	8 per system	16 per system
128-Port Asynchronous PCI	Primary drawer - slots 1, 3, 4, 11, 12, 13, 14, 10, 5, 9 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 14, 2, 6, 9, 13, 4, 7	8 per system	16 per system	32 per system
PCI SCSI-2 F/W Single-Ended PCI Single-Ended Ultra SCSI	Primary drawer - slots 1, 10, 14, 3, 9, 4, 12, 11, 5, 6, 2 Secondary drawer - slots 9, 2, 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14	40 per system (see note 4)	40 per system (see note 4)	40 per system (see note 4)
PCI SCSI-2 Fast/Wide Differential PCI Differential Ultra SCSI PCI Universal Differential Ultra SCSI	Primary drawer - slots 1, 10, 14, 3, 13, 4, 12, 11, 5, 6, 7 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14	40 per system (see note 4)	40 per system (see note 4)	40 per system (see note 4)
Eicon ISDN DIVA Pro PCI S/T	Primary drawer - slots 10, 14, 3, 13, 4, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	1 per system	1 per system	1 per system
Token ring PCI	Primary drawer - slots 10, 14, 3, 13, 4, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	9 per system	9 per system	9 per system
2-Port Multiprotocol PCI	Primary drawer - slots 10, 14, 3, 13, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	18 per system	18 per system	18 per system
10/100 Mbps PCI Fast EtherLink XL	Primary drawer - slots 10, 14, 3, 13, 4, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	12 per system (see note 6)	12 per system (see note 6)	12 per system (see note 6)

Adapter	Slot Usage	S70 Max	S7A Max	S80 Max
PCI Auto LANstreamer token ring	Primary drawer - slots 10, 14, 3, 13, 4, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	24 per system	24 per system	24 per system
Ethernet 10base2 PCI ¹⁰	Primary drawer - slots 10, 14, 3, 13, 4, 12, 11, 5, 9, 1, 2, 5 Secondary drawer - slots 1, 5, 10, 11, 3, 6, 12, 4, 7, 13, 8, 14, 9, 2	24 per system	24 per system	24 per system

Notes:

1. The use of the PCI SSA Multi-Initiator/RAID EL in the S70 I/O drawer limits the system usage to a 28°C (82°F) environment maximum. If installing a PCI SSA Multi-Initiator/RAID EL adapter, remove the screws from the blue plastic adapter guide and remove the guide before you install it in your S70 (save the guide and screws if you plan to install this adapter in a different system later). For maximum system performance, use a maximum of 16 adapters per system.

Use one adapter per bus and four adapters per drawer.

2. Remove the handle from your SSA adapter after you have installed it.
3. The manufacturer strongly recommends you locate the POWER GXT120P and GXT130P adapters in the primary I/O drawer. This placement provides you with the maximum amount of diagnostic feedback if your system encounters errors.

Do not use the POWER GXT130P in I/O drawers 3 or 4.

4. When using a non-Ultra2/LVD SCSI adapter, use a maximum of 40 storage adapters per system and 10 storage adapters per I/O drawer. See "Performance Limits" for more performance-related system maximum information.
5. Contact your marketing representative for information about this adapter. For performance-related information, see "Performance Limits".
6. See D.17.2, "Performance Limits" on page 706 for more performance-related system maximum information.
7. For optimum system performance, the 10/100 Mbps Ethernet PCI Adapter or 10/100 Mbps PCI Fast EtherLink XL adapters are recommended instead of this adapter.

8. For optimum system performance, models S70, S7A, and S80 support any combination of the following adapters, but the maximum combination varies according to your system. See “Performance Limits” for more information:
 - 155 TURBOWAYS ATM PCI MMF, when used in LAN emulation (LANE) mode
 - 155 TURBOWAYS ATM PCI UTP, when used in LAN emulation (LANE) mode
 - 10/100 Mbps Ethernet PCI, when used in 100 Mbps mode
9. See D.17.2, “Performance Limits” on page 706 for performance-related information. Since only three adapters are recommended per bus, limit placement to 11 adapters per drawer for optimum performance.
10. For optimum system performance, install a maximum of two PCI Dual Channel Ultra2 SCSI adapters per bus, seven per drawer, and 26 per system.
11. Do not install a TURBOWAYS 155 PCI MMF ATM or 10/100 Mbps Ethernet PCI Adapter in the slot to the right (looking from the rear) of the S/390 ESCON Channel PCI Adapter.
12. Do not install a PCI Dual Channel Ultra2 SCSI Adapter in the same PCI bus as a POWER GXT 120P or POWER GXT 130P.
13. Do not install more than a combination of six IBM ARTIC960RxD Quad Digital Trunk PCIs and IBM ARTIC960RxF PCI Adapters in this system.

D.17.2 Performance Limits

This section contains system maximum suggestions for certain adapters. Apply these suggestions for optimum system performance and use the *PCI Adapter Placement Reference*, SA38-0538 for specific slot plugging guidelines. For more performance related information, see “System Performance”.

The quantity of installed processors affects the performance of installed adapters. The table below provides some guidelines that compare performance characteristics of certain adapters. These performance characteristics vary according to the number of installed processors in a system.

Table 209 on page 707 provides adapter maximum information for running up to 12 processors (12-way) in the S70, S7A, and S80 systems. It also provides

maximum adapter information for running up to 24 processors (24-way) in the S80 system only.

Table 209. S70, S7A, and S80 Performance Limits

Adapter	S70 Max 12-Way 125 MHz	S7A Max 12-Way 262 MHz	S80 Max 12-Way 450 MHz	S80 Max 24-Way 450 MHz
Gigabit Ethernet-SX PCI MTU 1500 MTU 9000 (Jumbo) See note 3	Not Tested Not Tested	2 4	6 7 See note 1	8 8 See note 1
10/100 Mbps Ethernet PCI 10 Mbs FDX 100 Mbs FDX See note 3	26 5	26 10	26 15	26 22
TURBOWAYS 155 PCI MMF/UTP ATM MTU 1500 MTU 9180 See note 3	5 10	10 14	15 22	18 26
FDDI SK-NET See note 2	6	12	20	26

Notes:

1. For maximum performance, install only 2 Gigabit Ethernet-SX PCI Adapters per I/O drawer, 1 per bus.
2. This includes the following adapters:
 - SysKonnnect SK-NET FDDI-UP SAS PCI
 - SysKonnnect SK-NET FDDI-LP DAS PCI
 - SysKonnnect SK-NET FDDI-LP SAS PCI
3. For maximum performance, limit these adapters to eight per drawer, two per bus.

D.17.3 Model S70, Model S70A, and S80 Configuration Details

The e following configuration details are specific to the S-Series servers:

- The recommended location for the boot device (SCSI or Network) and graphics accelerator is within the primary I/O drawer (Drawer 0). This configuration provides service personnel with the maximum amount of diagnostic information if your system encounters errors in the boot sequence.
- Consider placing the AIX rootvg volume group in the primary I/O drawer. This allows AIX to boot if other I/O drawers are found off-line during boot.
- The default Boot Drive is in the lowest location in the center bay six-pack of the primary SCSI I/O drawer. If a boot source other than the internal SCSI disk is configured, the supporting SCSI Adapter must also be in the primary I/O drawer.
- SCSI-2 disk bays in an I/O drawer are connected and driven by a single SCSI Adapter and are installed in slot 9 or slot 13. Ultra SCSI disk bays can each be driven from separate Ultra SCSI Adapters. In this configuration, slot 13 drives the default boot device, and slot 6 drives the second Ultra SCSI disk bay. These adapters are optional on secondary I/O drawers.
- SCSI-2 Media bays in an I/O drawer are connected and driven by a single SCSI Adapter and are installed in slot 2. Ultra systems have the media bay driven from Slot 7. This adapter is optional on secondary I/O drawers.
- The Service Processor must occupy slot 8 of the primary I/O drawer.
- The PCI SCSI-2 Fast/Wide RAID Adapter installed in the primary SCSI I/O drawer can only be connected to external devices. The SCSI RAID adapter installed in secondary I/O drawers may be connected to internal disk bays.
- When possible, it is suggested that you place the PCI SSA Multi-Initiator/RAID EL, the Advanced SerialRAID SSA, the Advanced SerialRAID SSA Plus, and the PCI SCSI-2 Fast/Wide RAID on their own buses.
- Maximum limitations exist on adapters and devices that are specific to the adapter or device and are not interaction limits with others. This information can be found in the product sales manual.
- I/O slot 9 does not support any long PCI Adapter with backside components.

D.18 9076 RS/6000 SP Systems

These systems are designed for service representatives to install adapters. For information about specific slot requirements for adapters that your service

representative may be installing, refer to *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280.

The S70, S7A, and S80 systems can function as attached SP servers within the RS/6000 SP environment operating under control of the Parallel Systems Support Programs (PSSP) for AIX. This interconnection can be accomplished using the SP System Attachment for PCI or through an Ethernet connection. Some I/O adapters available on these systems are not supported on the SP environment and must be removed. Refer to the RS/6000 SP 9076-550 sales manual for a list of currently supported adapters.

A minimum of one Ethernet adapter is required for a S70, S7A, or S80 system to function as an attached server within the IBM RS/6000 SP environment.

D.19 Digital Trunk PCI Adapter Placement Considerations

When installing a Digital Trunk Quad PCI or IBM ARTIC960RxD Quad Digital Trunk PCI Adapter in a system, some additional guidelines must be followed to assure optimum system performance. The guidelines are noted below:

- If a system has a PCI Auto LANStreamer token ring or a token ring PCI installed along with either a Digital Trunk Quad PCI or a IBM ARTIC960RxD Quad Digital Trunk PCI Adapter, the PCI Auto LANStreamer token ring or token-ring PCI Adapters must be installed in a lower slot number relative to the Digital Trunk Quad PCI or IBM ARTIC960RxD Quad Digital Trunk PCI Adapter.
- Whenever multiple Digital Trunk Quad PCI or a IBM ARTIC960RxD Quad Digital Trunk PCI Adapters are installed in the same system, they should be installed in adjacent slots.
- Digital Trunk Quad PCI or IBM ARTIC960RxD Quad Digital Trunk PCI Adapters installed in the same system with a POWER GXT110P or a POWER GXT120P should not be installed on the same PCI bus as the POWER GXT110P or a POWER GXT120P.
- Digital Trunk Quad PCI or IBM ARTIC960RxD Quad Digital Trunk PCI Adapters installed in the same system with an ISA adapter should not be installed on the same PCI bus as the ISA bridge.

Appendix E. Site and Hardware Planning Information

This appendix includes site and hardware planning information and physical planning diagrams for the following RS/6000 models and racks:

- 7043 43P Model 140
- 7043 43P Model 150
- 7044 44P Model 170
- 7043 43P Model 260
- 7044 44P Model 270
- 7046 Model B50
- 7025 Model F50
- 7025 Model F80
- 7026 Model H50
- 7026 Model H70
- 7026 Model H80
- 7026 Model M80
- 7017 Models S70 and S70 Advanced
- 7017 Model S80
- 7014 Model S00 Rack
- 7015 Model R00 Rack

E.1 7043 43P Model 140 Site and Hardware Planning Information

Table 210 summarizes the site and hardware planning information for the 7043 43P Model 140.

Table 210. 7043 43P Model 140 Site and Hardware Planning Information

Dimensions	Desktop	Deskside
Height	165 mm (6.5 in)	450 mm (17.7 in)
Width	420 mm (16.5 in)	165 mm (6.5 in)
Width with optional vertical stand		235 mm (9.25 in)
Depth	460 mm (18.0 in)	460 mm (18.0 in)
Weight		
Minimum Configuration	14.5 kg	32 lbs
Maximum Configuration	18.2 kg	40 lbs

Electrical				
Power source loading (typical in kVA)	0.2			
Power source loading (max. in kVA)	0.4			
Voltage range (V AC) - US & World Trade	100 to 127 or 200 to 240 (switchable)			
Voltage range (V AC) - Japan	100 to 127 or 200 to 240 (autoranging)			
Frequency (Hertz)	50 to 60			
Thermal output (typical)	425 BTU/hr			
Thermal output (maximum)	850 BTU/hr			
Power requirements (typical)	125 Watts			
Power requirements (maximum)	250 Watts			
Power factor - US & World Trade	0.6			
Power factor - Japan	0.98			
Inrush current ²	less than 70 Amps at 120 V AC and at 240 V AC			
Maximum altitude ³	2135 m (7000 ft.)			
Temperature Range	Operating		Non-Operating	
	16 to 32 degrees C (60 to 90 degrees F)		10 to 43 degrees C (50 to 110 degrees F)	
Humidity (Noncondensing)	Operating		Non-Operating	
	8 to 80%		8 to 80%	
Wet Bulb Requirements	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions¹	Operating		Idle	
L _{WAd}	5.3 bels		5.0 bels	
L _{pAm}	43 dBA		43 dBA	
<L _{pA} > _m	40 dBA		40 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	76 mm (3 in)	76 mm (3 in)	50 mm (2 in)
Service	Install so that it can be taken to an area providing 457 mm (18 in) on the front and 457 mm (18 in) on the left side.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p>				

Figure 110 shows the physical planning diagram for the 7043 43P Model 140.




Figure 110. 7043 43P Model 140 Physical Planning Diagram

E.2 7043 43P Model 150 Site and Hardware Planning Information

Table 211 summarizes the site and hardware planning information for the 7043 43P Model 150.


Table 211. 7043 43P Model 150 Site and Hardware Planning Information

Dimensions	Desktop	Deskside
Height	165 mm (6.5 in)	450 mm (17.7 in)
Width	420 mm (16.5 in)	165 mm (6.5 in)
Width with optional vertical stand		235 mm (9.25 in)
Depth	460 mm (18.0 in)	460 mm (18.0 in)
Weight		
Minimum Configuration	14.5 kg	32 lbs
Maximum Configuration	18.2 kg	40 lbs
Electrical		
Power source loading (typical in kVA)		0.2
Power source loading (max. in kVA)		0.4
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	425 BTU/hr	
Thermal output (maximum)	850 BTU/hr	
Power requirements (typical)	125 Watts	
Power requirements (maximum)	250 Watts	
Power factor - US, World Trade/ Japan	0.98	
Inrush current ²	less than 70 Amps at 120 V AC and at 240 V AC	
Maximum altitude ³	2135 m (7000 ft.)	

Temperature Range	Operating 16 to 32 degrees C (60 to 90 degrees F)		Non-Operating 10 to 43 degrees C (50 to 110 degrees F)	
Humidity (Noncondensing)	Operating 8 to 80%		Non-Operating 8 to 80%	
Wet Bulb Requirements	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions¹	Operating		Idle	
L _{WA} d	5.4 bels		5.0 bels	
L _{pA} m	43 dBA		43 dBA	
<L _{pA} > _m	40 dBA		40 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	76 mm (3 in)	50 mm (2 in)	50 mm (2 in)
Service	Install so that it can be taken to an area providing 457 mm (18 in.) on the front and 457 mm (18 in) on the left side.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions. 2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints. 3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

Figure 111 shows the physical planning diagram for the 7043 43P Model 150.

Figure 111. 7043 43P Model 150 Physical Planning Diagram



E.3 7044 44P Model 170 Site and Hardware Planning Information

Table 212 summarizes the site and hardware planning information for the 7044 44P Model 170.

Table 212. 7044 44P Model 170 Site and Hardware Planning Information

Dimensions	
Height	490 mm (19.25 in)
Width	200 mm (7.9 in)
Width ⁴	235 mm (9.25 in)
Depth	515 mm (20.25 in)

Weight				
Minimum Configuration	17.7 kg		39 lbs	
Maximum Configuration	20.4 kg		45 lbs	
Electrical				
Power source loading (typical in kVA)		0.23		
Power source loading (max. in kVA)		0.4		
Voltage range (V ac)				
- US, World Trade, and Japan		100 to 127 or 200 to 240 (autoranging)		
Frequency (Hertz)		50 to 60		
Thermal output (typical)		752 BTU/hr		
Thermal output (maximum)		1368 BTU/hr		
Power requirements (typical)		220 Watts		
Power requirements (maximum)		400 Watts		
Power factor - US, World Trade, Japan		0.98		
Inrush current ³		less than 60 amps at 120 V ac and at 240 V ac		
Maximum altitude		2135 m (7000 ft.)		
Temperature Range				
	Operating		Non-Operating (Power Off)	
	16 to 32 degrees C		10 to 43 degrees C	
	(60 to 90 degrees F)		(50 to 110 degrees F)	
Humidity Requirements				
(Noncondensing)	Operating		Non-Operating (Power Off)	
	8 to 80%		8 to 80%	
Wet Bulb Requirements	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions¹				
	Operating		Idle	
L _{WA} d	5.5 bels		5.4 bels	
L _{pA} m	NA		NA	
<L _{pA} > _m	38 dBA		37 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances				
	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	76 mm (3 in)	0 mm (0 in)	0 mm (0 in)
Service				
Install so that it can be taken to an area providing 457 mm (18 in) on the front and 457 mm (18 in) on the left side.				
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p> <p>4. Width measurement With feet extended.</p>				

Figure 112 shows the physical planning diagram for the 7044 44P Model 170.




Figure 112. 7044 44P Model 170 Physical Planning Diagram

E.4 7043 43P Model 260 Site and Hardware Planning Information

Table 213 summarizes the site and hardware planning information for the 7043 43P Model 260.

Table 213. 7043 43P Model 260 Site and Hardware Planning Information

Dimensions		
Height	610 mm (24.0 in)	
Width	222 mm (8.7 in)	
Width with Pedestal	340 mm (13.4 in)	
Depth	713 mm (28.1 in)	
Weight		
Minimum Configuration	37 kg	80 lbs
Maximum Configuration	45 kg	97 lbs
Electrical		
Power source loading (max. in kVA)	0.41	
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	883 BTU/hr	
Thermal output (maximum)	1324 BTU/hr	
Power requirements (typical)	259 Watts	
Power requirements (maximum)	388 Watts	
Power factor	0.89 to 0.98	
Inrush current ²	16 Amps at 120 V AC, 21Amps at 240 V AC	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	16 to 32 degrees C (60 to 90 degrees F)	10 to 43 degrees C (50 to 110 degrees F)

	Operating		Non-Operating	
Humidity (Noncondensing)	8 to 80%		8 to 80%	
Wet Bulb Requirements	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions¹	Operating		Idle	
L _{WAd}	5.5 bels		5.4 bels	
L _{pAm}	NA		NA	
<L _{pA} > _m	37 dBA		36 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	152 mm (6 in)	51 mm (2 in)	51 mm (2 in)
Service	Install so that it can be taken to an area providing 457 mm (18 in) on the front and 457 mm (18 in) on the left side.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions. 2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints. 3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

Figure 113 shows the physical planning diagram for the 7043 43P Model 260.

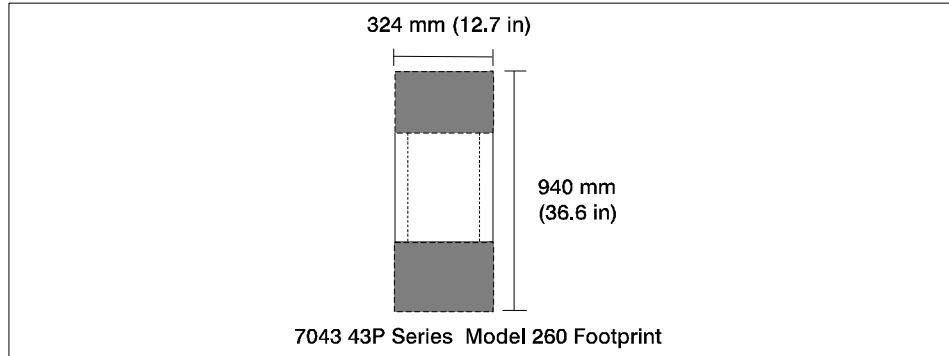


Figure 113. 7043 43P Model 260 Physical Planning Diagram

E.5 7044 44P Model 270 Site and Hardware Planning Information

Table 214 summarizes the site and hardware planning information for the 7044 44P Model 270.

Table 214. 7044 44P Model 270 Site and Hardware Planning Information

Dimensions					
Height with Pedestal	610 mm (24.0 in)				
Width with Pedestal	340 mm (13.4 in)				
Depth with Pedestal	713 mm (28.1 in)				
Weight					
Minimum Configuration	37 kg	80 lbs			
Maximum Configuration	45 kg	97 lbs			
Electrical					
Power source loading (max. in kVA)	0.47				
Voltage range (V ac)	100 to 127 or 200 to 240 (autoranging)				
Frequency (Hertz)	50 to 60				
Thermal output (typical)	1012 BTU/hr				
Thermal output (maximum)	1518 BTU/hr				
Power requirements (typical)	297 Watts				
Power requirements (maximum)	445 Watts				
Power factor	0.92 to 0.99				
Inrush current ³	30 amps at 120 V ac, 32 amps at 240 V ac				
Maximum altitude	2135 m (7000 ft.)				
Temperature Range		Operating	Non-Operating		
		16 to 32 degrees C (60 to 90 degrees F)	10 to 43 degrees C (50 to 110 degrees F)		
Humidity Requirements (Noncondensing)		Operating	Non-Operating		
		8 to 80%	8 to 80%		
Wet Bulb Requirements		23 degrees C (73 degrees F)	27 degrees C (80 degrees F)		
Noise Emissions¹		Operating	Idle		
L _{WAd}		5.5 bels	5.4 bels		
L _{pAm}		NA	NA		
<L _{pA} > _m		37 dBA	36 dBA		
Impulsive or prominent discrete tones		No	No		
Clearances		Front	Back	Left	Right
Install/Air Flow²		76 mm (3 in)	152 mm (6 in)	51 mm (2 in)	51 mm (20 in)
Service		Install so that it can be taken to an area providing 457 mm (18 in) on the front and 457 mm (18 in) on the left side.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p>					

Figure 114 shows the physical planning diagram for the 7044 44P Model 270.




Figure 114. 7044 44P Model 270 Physical Planning Diagram

E.6 7046 Model B50 Site and Hardware Planning Information

Table 215 summarizes the site and hardware planning information for the 7046 Model B50.

Table 215. 7046 Model B50 Site and Hardware Planning Information

Dimensions	Individual	Including pull handles
Height	88 mm (3.5 in)	88 mm (3.5 in)
Width	447 mm (17.6 in)	447 mm (17.6 in)
Depth	612.1 mm (24.1 in)	751.8 mm (29.6 in)
Weight	14.5 kg (32 lb.)	14.5 kg (32 lb)
Electrical		
Power source loading (max. in kVA)		0.147
Voltage range (V AC)	10 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (maximum)	478 BTU/hr	
Power requirements (maximum)	140 Watts	
Power factor - US, World Trade/ Japan	0.95	
Inrush current ²	40 amps	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	10 to 40 degrees C (60 to 104 degrees F)	10 to 52 degrees C (50 to 126 degrees F)
Humidity (Noncondensing)	Operating	Non-Operating
	8 to 80%	8 to 80%
Wet Bulb Requirements	27 degrees C (80 degrees F)	27 degrees C (80 degrees F)
Noise Emissions¹	Operating	Idle
L _{WAd}	5.2 bels	5.0 bels
L _{pAm}	N/A	N/A

$\langle L_{pA} \rangle_m$	35 dBA		30 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	76 mm (3 in)	50 mm (2 in)	50 mm (2 in)
Service				
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions. 2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints. 3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

E.7 7025 Model F50 Site and Hardware Planning Information

Table 216 summarizes the site and hardware planning information for the system rack of the 7025 Model F50.

Table 216. 7025 Model F50 Site and Hardware Planning Information

Dimensions		
Height	620 mm	24.3 in
Width	245 mm	9.6 in
Width with Pedestal	350 mm	13.7 in
Depth	695 mm	27.3 in
Depth with Pedestal	745 mm	29.3 in
Weight		
Minimum Configuration	30 kg	65 lbs
Maximum Configuration	55 kg	120 lbs
Electrical		
Power source loading (typical in kVA)	0.52	
Power source loading (max. in kVA)	0.56	
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	975 BTU/hr	
Thermal output (maximum)	2050 BTU/hr	
Power requirements (typical)	285 Watts	
Power requirements (maximum)	600 Watts	
Power factor	0.8 - 0.96	
Inrush current ³	50 Amps	
Maximum altitude	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	16 to 32 degrees C (60 to 90 degrees F)	10 to 43 degrees C (50 to 110 degrees F)
Humidity (Noncondensing)	Operating	Non-Operating
	8 to 80%	8 to 80%
Wet Bulb Requirements	23 degrees C (73 degrees F)	27 degrees C (80 degrees F)
Noise Emissions¹	Operating	Idle
L_{WAd}	5.8 bels	5.5 bels

L_{pAm}	N/A		N/A	
$\langle L_{pA} \rangle_m$	41 dBA		38 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow²	76 mm (3 in)	152 mm (6 in)	51 mm (2 in)	51 mm (2 in)
Service	Install so that it can be moved to an area providing 457 mm (18 in) on the front and 457 mm (18 in) on the left side.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions. 2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints. 3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

Figure 115 shows the physical planning diagram for the 7025 Model F50.




Figure 115. 7025 F50 Physical Planning Diagram

E.8 7025 Model F80 Site and Hardware Planning Information

Table 217 summarizes the site and hardware planning information for the 7025 Model F80.

Table 217. 7025 Model F80 Site and Hardware Planning Information

Dimensions		
Height	610 mm	24.0 in
Width	483 mm	19.0 in
Depth	728 mm	28.7 in
Weight		
Minimum Configuration	70 kg	155 lbs
Maximum Configuration	95 kg	209 lbs

Electrical			
Power source loading (typical in kVA)	0.59		
Power source loading (max. in kVA)	0.86		
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)		
Frequency (Hertz)	50 to 60		
Thermal output (typical)	1920 BTU/hr		
Thermal output (maximum)	2867 BTU/hr		
Power requirements (typical)	560 Watts		
Power requirements (maximum)	840 Watts		
Power factor	0.95		
Inrush current ²	70 Amps		
Maximum altitude ³	2135 m (7000 ft.)		
Temperature Range	Operating⁴		Non-Operating
	10 to 38 degrees C (50 to 100 degrees F)		10 to 43 degrees C (50 to 110 degrees F)
	Operating		Non-Operating
Humidity (Noncondensing)	8 to 80%		8 to 80%
Wet Bulb Requirements	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)
Noise Emissions¹	Operating		Idle
L _{WAd}	6.1 bels		5.9 bels
L _{pAm}	N/A		N/A
<L _{pA} > _m	43 dBA		40 dBA
Impulsive or prominent discrete tones	No		No
Clearances	Front	Back	Left Right
Install/Air Flow	76 mm (3 in)	152 mm (6 in)	51 mm (2 in) 51 mm (2 in)
Service	Install so that it can be moved to an area providing 457 mm (18 in.) on the front and 457 mm (18 in) on the left side.		
<p>1. See -- Heading 'EMISS' unknown -- for definitions of noise emissions positions.</p> <p>2. The amount of space needed by the unit during normal operation is indicated by broken lines on the footprints.</p> <p>3. Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.</p> <p>4. A) Dry bulb temperature derating at high altitude: Upper limit temperature must be derated 1.0 deg C per 137 m of elevation beyond 1295 m above sea level. (1 deg F per 250 ft above 4250 feet).</p> <p>B) Wet bulb temperature derating at high altitude: Upper limit temperature must be derated 1 deg C per 274 m of elevation beyond 1372m above sea level (1 deg F per 500 ft above 4500 feet).</p>			

Figure 116 on page 723 shows the physical planning diagram for the 7025 Model F80.




Figure 116. 7025 F80 Physical Planning Diagram

E.9 7026 Model H70 Site and Hardware Planning Information

Table 218 summarizes the site and hardware planning information for the 7026 Model H70.

Table 218. 7026 Model H70 Site and Hardware Planning Information

Dimensions		
Height	350 mm	13.8 in
Width	443 mm	17.4 in
Depth	875 mm	34.2 in
Weight		
Minimum Configuration	71 kg	157 lbs
Maximum Configuration	89 kg	195 lbs
Electrical		
Power source loading (typical in kVA)	0.46	
Power source loading (max. in kVA)	0.691	
Voltage range (V AC)	200 to 240 V AC	
Frequency (Hertz)	50 to 60 Hz	
Thermal output (typical)	1485 BTU/hr	
Thermal output (maximum)	2818 BTU/hr	
Power requirements (typical)	434 Watts	
Power requirements (maximum)	650 Watts	
Power factor	0.9-0.98	
Inrush current ²	50 Amps	
Maximum altitude ³	915 m (3000 ft.)	

Temperature Range	Operating 10 to 40 C (50 to 104 F)		Non-Operating Shipping 10 to 52 C (34 to 125.6 F)	
Humidity (Noncondensing)	Operating 8 to 80%		Non-Operating 8 to 80%	
Wet Bulb Requirements	27 C (80.6F)		27 C (80.6 F)	
Noise Emissions¹				
L _{WAd}	6.0 bels		5.9 bels	
L _{pAm}	NA		NA	
<L _{pA} > _m	43 dBA		40 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Service	1650 mm (65 in)	1015 mm (40 in)	915 mm (36 in)	915 mm (36 in)
Install/Air Flow	Maintenance of proper service clearances should allow proper air flow.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions. 2. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle. 3. For altitudes above 915 meters, the maximum temperature limit is derated by 1 degree C for every 137 meters of elevation above 915 meters.				

The 7026 Model H70 can be installed in a Model S00 or Model R00 rack. See Figure 119 on page 739 for the Model S00 rack physical planning diagram or Figure 120 on page 740 for the Model R00 rack physical planning diagram.

E.10 7026 Model H80 CEC Drawer Site and Hardware Planning Info.

The Model H80 includes two drawers. They are the Central Electronics Complex (CEC) drawer with an I/O Drawer. For technical information on the I/O Drawer, see Table 220 on page 726 for site and hardware planning information.

Table 219 summarizes the site and hardware planning information for the 7026 Model H80 CEC Drawer.

Table 219. 7026 Model H80 CEC Drawer Site and Hardware Planning Info

Dimensions		
Height	218 mm	8.58 in 5 (EIA Units)
Width	445 mm	17.5 in
Depth	820 mm	32.3 in
Weight		
Minimum Configuration	41.0 kg	90 lbs
Maximum Configuration	52.0 kg	115 lbs
Electrical		
Power source loading (typical in kVA)		0.32

Power source loading (max. in kVA)	0.48			
Voltage range (V AC)	200 to 240 V AC			
Frequency (Hertz)	50 to 60 Hz			
Thermal output (typical)	1025 BTU/hr			
Thermal output (maximum)	1536 BTU/hr			
Power requirements (typical)	300 Watts			
Power requirements (maximum)	450 Watts			
Power factor	0.95			
Inrush current ²	40 Amps			
Maximum altitude ³	2135 m (7000 ft.)			
Temperature Range	Operating		Non-Operating	
	10 to 40 C (50 to 104 F)		10 to 52 C (50 to 125.6 F)	
Humidity (Noncondensing)	Operating		Non-Operating	
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		8 to 80%	
Wet Bulb Requirements				
Without tape drive	27 C (80.6 F)		27 C (80.6 F)	
With tape drive	23 C (73 F)		27 C (80.6 F)	
Noise Emissions¹				
L _{WA} d	5.8 bels		5.8 bels	
L _{pA} m	NA		NA	
<L _{pA} > _m	45 dBA		45 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of proper service clearances should allow proper air flow.			
Service	(See service clearances for the 7014 T00 Rack)			
<p>1. Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.</p> <p>2. For altitudes above 915 meters, the maximum temperature limit is derated by 1 degree C for every 137 meters of elevation above 915 meters.</p> <p>3. See -- Heading 'EMISS' unknown -- for definitions of noise emissions positions.</p>				

Table 220. Site and Hardware Planning Information for I/O Drawer

Dimensions				
Height	218 mm		8.6 in	
Width	445 mm		17.5 in	
Depth	820 mm		32.3 in	
Weight				
Minimum Configuration	41 kg		90 lbs	
Maximum Configuration	52 kg		115 lbs	
Electrical				
Power source loading (typical in kVA)		0.23		
Power source loading (max. in kVA)		0.54		
Voltage range (V AC)		200 to 240 V AC		
Frequency (Hertz)		50 to 60 Hz		
Thermal output (typical)		750 BTU/hr		
Thermal output (maximum)		1095 BTU/hr		
Power requirements (typical)		220 Watts		
Power requirements (maximum)		515 Watts		
Power factor		0.95		
Inrush current ²		41 Amps		
Maximum altitude ³		2135 m (7000 ft.)		
Temperature Range	Operating		Non-Operating	
	10 to 40 C (50 to 104 F)		10 to 52 C (50 to 125.6 F)	
	Operating		Non-Operating	
Humidity (Noncondensing)				
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		20 to 80%	
Wet Bulb Requirements				
Without tape drive	27 C (80 F)		27 C (80 F)	
With tape drive	23 C (80 F)		27 C (80 F)	
Noise Emissions¹				
L _{WA} d	5.8 bels		5.8 bels	
L _{pA} m	NA		NA	
<L _{pA} > _m	45 dBA		45 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow	1650 mm (65 in)	1015 mm (40 in)	915 mm (36 in)	915 mm (36 in)
Service	Maintenance of proper service clearances should allow proper air flow.			
1. See "Noise Emission Notes" on page 2 for definitions of noise emissions positions.				
2. Noise emissions data is based on the following configuration: the drawer is mounted in a T00 Rack and a power distribution unit.				
3. Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.				

E.11 7026 Model M80 CEC Drawer Site and Hardware Planning Info.

The Model M80 includes two drawers. They are the Central Electronics Complex (CEC) drawer with an I/O Drawer. For technical information on the I/O Drawer, see Table 222 on page 728 for site and hardware planning information.

Table 221 summarizes the site and hardware planning information for the 7026 Model M80 CEC Drawer.

Table 221. 7026 Model M80 CEC Drawer Site and Hardware Planning Info

Dimensions		
Height	355.6 mm	14.0 in
Width	445.5 mm	17.5 in
Depth	825.5 mm	32.5 in
Weight		
Minimum Configuration	69.7 kg	158 lbs
Maximum Configuration	74.6 kg	169 lbs
Electrical		
Power source loading (typical in kVA)	0.39	
Power source loading (max. in kVA)	0.6	
Voltage range (V AC)	200 to 240 V AC	
Frequency (Hertz)	50 to 60 Hz	
Thermal output (typical)	1265 BTU/hr	
Thermal output (maximum)	1877 BTU/hr	
Power requirements (typical)	370 Watts	
Power requirements (maximum)	550 Watts	
Power factor	0.95	
Inrush current ²	34 Amps	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	10 to 40 C (50 to 104 F)	10 to 40 C (52 to 125 F)
	Operating	Non-Operating
Humidity (Noncondensing)		
Without tape drive	8 to 80%	8 to 80%
With tape drive	20 to 80%	8 to 80%
Wet Bulb Requirements		
Without tape drive	27 C (80 F)	27 C (80 F)
With tape drive	27 C (80 F)	27 C (80 F)
Noise Emissions¹		
With M80 CEC Drawer only		
L _{WAd}	6.4 bels	6.4 bels
L _{pAm}	NA	NA

<L _{pA} > _m	48 dBA	48 dBA		
Impulsive or prominent discrete tones	No	No		
Noise Emissions¹				
With M80 CEC drawer & Primay I/O Drawer				
L _{WAd}	6.5 bels	6.5 bels		
L _{pAm}	NA	NA		
<L _{pA} > _m	49 dBA	49 dBA		
Impulsive or prominent discrete tones	No	No		
Clearances	Front	Back	Left	Right
Install/Air Flow	1650 mm (65 in)	1015 mm (40 in)	915 mm (36 in)	915 mm (36 in)
Service	Maintenance of proper service clearances should allow proper air flow.			
1. See -- Heading 'EMISS' unknown -- for definitions of noise emissions positions.				
2. Noise emissions data is based on the following configuration: the drawer is mounted in a T00 Rack and a power distribution unit is installed in the rack and the system is operating in a normal environment of 25 °C (78 °F)				
3. Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.				

Table 222. Site and Hardware Planning Information for I/O Drawer

Dimensions		
Height	218 mm	8.6 in
Width	445 mm	17.5 in
Depth	820 mm	32.3 in
Weight		
Minimum Configuration	41 kg	90 lbs
Maximum Configuration	52 kg	115 lbs
Electrical		
Power source loading (typical in kVA)	0.23	
Power source loading (max. in kVA)	0.34	
Voltage range (V AC)	200 to 240 V AC	
Frequency (Hertz)	50 to 60 Hz	
Thermal output (typical)	750 BTU/hr	
Thermal output (maximum)	1095 BTU/hr	
Power requirements (typical)	220 Watts	
Power requirements (maximum)	512 Watts	
Power factor	0.95	
Inrush current ²	41 Amps	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	10 to 40 C (50 to 104 F)	10 to 52 C (50 to 125.6 F)
Humidity (Noncondensing)	Operating	Non-Operating
Without tape drive	8 to 80%	8 to 80%
With tape drive	20 to 80%	20 to 80%

Wet Bulb Requirements				
Without tape drive	27 C (80 F)		27 C (80 F)	
With tape drive	23 C (80 F)		27 C (80 F)	
Noise Emissions¹				
L _{WAd}	5.8 bels		5.8 bels	
L _{pAm}	NA		NA	
<L _{pA} > _m	45 dBA		45 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances	Front	Back	Left	Right
Install/Air Flow	1650 mm (65 in)	1015 mm (40 in)	915 mm (36 in)	915 mm (36 in)
Service	Maintenance of proper service clearances should allow proper air flow.			
1. See "Noise Emission Notes" on page 2 for definitions of noise emissions positions.				
2. Noise emissions data is based on the following configuration: the drawer is mounted in a T00 Rack and a power distribution unit.				
3. Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.				

E.12 7017 Models S70 and S70 Advanced Site and Hardware Planning

In this section, site and hardware planning information is covered for the following:

- Models S70 and S70 Advanced system rack
- Model S70 SCSI I/O drawer (7 EIA)
- Model S70 Advanced I/O drawer (10 EIA)
- Models S70 and S70 Advanced I/O rack

In addition, physical planning diagrams are included for the following:

- Models S70 and S70 Advanced rack (AC Systems)

Table 223 summarizes the site and hardware planning information for the system rack of the 7017 Models S70 and S70 Advanced.

Table 223. 7017 Models S70 and S7A Rack Physical Planning Information

Dimensions				
Height	1577 mm	62.0 in		
Width	567 mm	22.3 in		
Depth	1041 mm	40.9 in		
Weight				
Minimum (Configuration Dependent)	400 kg	880 lbs		
Electrical				
Power source loading (maximum in kVA)	1.887 kVA			
Voltage range (V AC)	200 to 240			
Frequency (Hertz)	50 to 60			
Thermal output (Maximum)	5796 BTU/hr.			
Power requirements (Maximum)	1698 Watts			
Power factor	0.9			
Inrush current ³	102 Amps			
Maximum altitude	2135 m (7000 ft.)			
Temperature Range^{4,5}				
	Operating	Non-Operating		
	10 to 27.8 degrees C (50 to 100 degrees F)	1 to 60 degrees C (34 to 140 degrees F)		
Humidity (Noncondensing)				
	Operating	Non-Operating		
	8 to 80%	8 to 80%		
Wet Bulb Requirements⁶				
	23 degrees C (73 degrees F)	23 degrees C (73 degrees F)		
Noise Emissions^{1,2}				
	Operating	Idle		
L _{WAd}	7.0 bels	7.0 bels		
L _{pAm}	N/A	N/A		
<L _{pA} > _m	N/A	N/A		
Impulsive or prominent discrete tones	No	No		
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagram for rack.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. Noise emissions data for Models S70 and S7A are based on a system with the doors closed.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p> <p>4. The use of the PCA SSA Multi-Initiator/RAID EL in the Model S70 I/O Drawer limits the system usage to a 28 degrees C (82 degrees F) environment maximum.</p> <p>5. The upper limit of the dry bulb temperature must be derated 1 degree C per 137 m (450 ft.) above 1295 m (4250 ft.).</p> <p>6. The upper limit of the wet bulb temperature must be derated 1 degree C per 274 m (882 ft.) elevation above 1370 m (4500 ft.).</p>				

Table 224 summarizes the site and hardware planning information for the SCSI I/O drawer for the 7017 Model S70.

Table 224. 7017 Model S70 SCSI I/O Drawer 7 EIA Planning Information

Dimensions				
Height	306.2 mm		12.1 in	
Width	442.4 mm		17.4 in	
Depth	748.2 mm		29.5 in	
Weight				
Minimum configuration	43 kg		95 lbs	
Maximum configuration	61 kg		135 lbs	
Electrical				
	AC		DC	
Power source loading (typical in kVA)	0.4		0.4	
Power source loading (max. in kVA)	1.0		1.0	
Voltage range	200 to 240 V AC		40 to 60 V DC	
Frequency (Hertz)	50 to 60		N/A	
Thermal output (typical)	1228 BTU/hr		1365 BTU/hr	
Thermal output (maximum)	3071 BTU/hr		3412 BTU/hr	
Power requirements (typical)	360 Watts		400 Watts	
Power requirements (maximum)	900 Watts		1000 Watts	
Power factor	0.9		N/A	
Inrush current ³	120 Amps		300 Amps	
Maximum altitude	2135 m (7000 ft.)		2135 m (7000 ft.)	
Temperature Range⁴				
	Operating		Non-Operating	
	10 to 40 degrees C (50 to 104 degrees F)		10 to 52 degrees C (50 to 125.6 degrees F)	
Humidity (Noncondensing)				
	Operating		Non-Operating	
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		20 to 80%	
Wet Bulb Requirements				
Without tape drive	27 degrees C (80 degrees F)		27 degrees C (80 degrees F)	
With tape drive	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions^{1,2}				
	Operating		Idle	
L _{WAd}	5.9 bels		5.8 bels	
L _{pAm}	N/A		N/A	
<L _{pA} > _m	39 dBA		38 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. Noise emissions data for Models S70 and S7A are based on a system with the doors closed.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p> <p>4. The use of the PCA SSA Multi-Initiator/RAID EL in the Model S70 I/O Drawer limits the system usage to a 28 degrees C (82 degrees F) environment maximum.</p>				

Table 225 summarizes the site and hardware planning information for the SCSI I/O drawer of the 7017 Model S70 Advanced.

Table 225. Model S70 Advanced I/O Drawer 10 EIA Planning Information

Dimensions				
Height	440.0 mm		17.3 in	
Width	443.2 mm		17.5 in	
Depth	843.2 mm		33.2 in	
Weight				
Minimum configuration	89 kg		195 lbs	
Maximum configuration	93 kg		205 lbs	
Electrical				
		AC		
Power source loading (typical in kVA)		0.4		
Power source loading (max. in kVA)		1.0		
Voltage range		200 to 240 V AC		
Frequency (Hertz)		50 to 60		
Thermal output (typical)		1228 BTU/hr		
Thermal output (maximum)		3071 BTU/hr		
Power requirements (typical)		360 Watts		
Power requirements (maximum)		900 Watts		
Power factor		0.96		
Inrush current ³		170 Amps		
Maximum altitude		2135 m (7000 ft.)		
Temperature Range				
	Operating		Non-Operating	
	10 to 40 degrees C (50 to 104 degrees F)		10 to 52 degrees C (50 to 125.6 degrees F)	
Humidity (Noncondensing)				
	Operating		Non-Operating	
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		20 to 80%	
Wet Bulb Requirements				
Without tape drive	27 degrees C (80 degrees F)		27 degrees C (80 degrees F)	
With tape drive	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions^{1,2}				
	Operating		Idle	
L _{WAd}	5.9 bels		5.3 bels	
L _{pAm}	N/A		N/A	
<L _{pA} > _m	N/A		N/A	
Impulsive or prominent discrete tones	No		No	
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.				
2. Noise emissions data for Models S70 and S7A are based on the I/O drawer mounted in a rack. See S70 and S7A I/O Rack.				
3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

Table 226 summarizes the site and hardware planning information for the I/O rack of the 7017 Models S70 and S70 Advanced.

Table 226. 7017 Models S70 and S70 Advanced I/O Rack Planning Information

Dimensions				
Height		1577 mm		62.0 in
Width		650 mm		25.5 in
Depth		1019 mm		40.1 in
Weight ¹ (Base Rack)		159 kg		349 lbs
Electrical	(see specifications for drawers or enclosures)			
Temperature Range	(see specifications for drawers or enclosures)			
Humidity (Noncondensing)	(see specifications for drawers or enclosures)			
Wet Bulb Requirements	(see specifications for drawers or enclosures)			
Noise Emissions	(see specifications for drawers or enclosures)			
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
1. Configuration-dependent, base weight plus weight of drawers				

Figure 117 shows the physical planning diagram for the 7017 Models S70 and S70 Advanced Rack for AC systems.




Figure 117. 7017 Models S70 and S7A Rack Physical Planning Diagram

E.13 7017 Model S80 Site and Hardware Planning Information

In this section, site and hardware planning information is covered for the following:

- Models S80

In addition, physical planning diagrams are included for the following:

- Models S80 rack (AC Systems)

Table 227 summarizes the site and hardware planning information for the system rack of the 7017 Model S80.

Table 227. 7017 Model S80 CEC Rack Physical Planning Information

Dimensions				
Height	1577 mm	62.0 in		
Width	567 mm	22.3 in		
Depth	1041 mm	40.9 in		
Weight				
Minimum (Configuration Dependent)	400 kg	880 lbs		
Electrical				
Power source loading (maximum in kVA)	2.129 kVA			
Voltage range (V AC)	200 to 245			
Frequency (Hertz)	50 to 60			
Thermal output (Maximum)	6904 BTU/hr.			
Power requirements (Maximum)	2023 Watts			
Power factor	0.92 to 0.98			
Inrush current ³	43 amps			
Maximum altitude	2135 m (7000 ft.)			
Temperature Range^{4,5}				
	Operating	Non-Operating		
	10 to 37.8 degrees C (50 to 100 degrees F)	1 to 60 degrees C (34 to 140 degrees F)		
Humidity (Noncondensing)				
	Operating	Non-Operating		
	8 to 80%	8 to 80%		
Wet Bulb Requirements⁶				
	23 degrees C (73 degrees F)	23 degrees C (73 degrees F)		
Noise Emissions^{1,2}				
	Operating	Idle		
L _{WAd}	7.0 bels	7.0 bels		
L _{pAm}	N/A	N/A		
<L _{pA} > _m	N/A	N/A		
Impulsive or prominent discrete tones	No	No		
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagram for rack.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. Noise emissions data for Models S70 and S7A are based on a system with the doors closed.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p> <p>4. The use of the PCA SSA Multi-Initiator/RAID EL in the Model S70 I/O Drawer limits the system usage to a 28 degrees C (82 degrees F) environment maximum.</p> <p>5. The upper limit of the dry bulb temperature must be derated 1 degree C per 137 m (450 ft.) above 1295 m (4250 ft.).</p> <p>6. The upper limit of the wet bulb temperature must be derated 1 degree C per 274 m (882 ft.) elevation above 1370 m (4500 ft.).</p>				

Table 228 summarizes the site and hardware planning information for the SCSI I/O drawer for the 7017 Model S80.

Table 228. 7017 Model S80 SCSI I/O Drawer 7 EIA Planning Information

Dimensions				
Height	440.0 mm			17.3 in
Width	443.2 mm			17.5 in
Depth	843.2 mm			33.2 in
Weight				
Minimum configuration	89 kg			195 lbs
Maximum configuration	93 kg			205 lbs
Electrical				
Power source loading (typical in kVA)				0.4
Power source loading (max. in kVA)				1.0
Voltage range				200 to 240 V AC
Frequency (Hertz)				50 / 60
Thermal output (typical)				1228 BTU/hr
Thermal output (maximum)				3071 BTU/hr
Power requirements (typical)				360 watts
Power requirements (maximum)				900 watts
Power factor				0.96
Inrush current ³				170 amps
Maximum altitude				2135 m (7000 ft.)
Temperature Range⁴				
	Operating		Non-Operating	
	10 to 40 degrees C (50 to 104 degrees F)		10 to 52 degrees C (50 to 125.6 degrees F)	
Humidity (Noncondensing)				
	Operating		Non-Operating	
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		20 to 80%	
Wet Bulb Requirements				
Without tape drive	27 degrees C (80 degrees F)		27 degrees C (80 degrees F)	
With tape drive	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions^{1,2}				
	Operating		Idle	
L _{WAd}	5.9 bels		5.8 bels	
L _{pAm}	N/A		N/A	
<L _{pA} > _m	39 dBA		38 dBA	
Impulsive or prominent discrete tones	No		No	
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
<p>1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.</p> <p>2. Noise emissions data for Models S70 and S7A are based on a system with the doors closed.</p> <p>3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.</p> <p>4. The use of the PCA SSA Multi-Initiator/RAID EL in the Model S70 I/O Drawer limits the system usage to a 28 degrees C (82 degrees F) environment maximum.</p>				

Table 229 summarizes the site and hardware planning information for the SCSI I/O drawer of the 7017 Model S70 Advanced.

Table 229. Model S70 Advanced I/O Drawer 10 EIA Planning Information

Dimensions				
Height	440.0 mm		17.3 in	
Width	443.2 mm		17.5 in	
Depth	843.2 mm		33.2 in	
Weight				
Minimum configuration	89 kg		195 lbs	
Maximum configuration	93 kg		205 lbs	
Electrical				
		AC		
Power source loading (typical in kVA)		0.4		
Power source loading (max. in kVA)		1.0		
Voltage range		200 to 240 V AC		
Frequency (Hertz)		50 to 60		
Thermal output (typical)		1228 BTU/hr		
Thermal output (maximum)		3071 BTU/hr		
Power requirements (typical)		360 Watts		
Power requirements (maximum)		900 Watts		
Power factor		0.96		
Inrush current ³		170 Amps		
Maximum altitude		2135 m (7000 ft.)		
Temperature Range				
	Operating		Non-Operating	
	10 to 40 degrees C (50 to 104 degrees F)		10 to 52 degrees C (50 to 125.6 degrees F)	
Humidity (Noncondensing)				
	Operating		Non-Operating	
Without tape drive	8 to 80%		8 to 80%	
With tape drive	20 to 80%		20 to 80%	
Wet Bulb Requirements				
Without tape drive	27 degrees C (80 degrees F)		27 degrees C (80 degrees F)	
With tape drive	23 degrees C (73 degrees F)		27 degrees C (80 degrees F)	
Noise Emissions^{1,2}				
	Operating		Idle	
L _{WAd}	5.9 bels		5.3 bels	
L _{pAm}	N/A		N/A	
<L _{pA} > _m	N/A		N/A	
Impulsive or prominent discrete tones	No		No	
Clearances				
	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
1. See "Noise Emission Notes" on page 742 for definitions of noise emissions positions.				
2. Noise emissions data for Models S70 and S7A are based on the I/O drawer mounted in a rack. See S70 and S7A I/O Rack.				
3. Inrush currents occur only at initial application of power; no inrush occurs during normal power off-on cycle.				

Table 230 summarizes the site and hardware planning information for the I/O rack of the 7017 Models S70 and S70 Advanced.

Table 230. 7017 Models S70 and S70 Advanced I/O Rack Planning Information

Dimensions				
Height		1577 mm		62.0 in
Width		650 mm		25.5 in
Depth		1019 mm		40.1 in
Weight ¹ (Base Rack)		159 kg		349 lbs
Electrical	(see specifications for drawers or enclosures)			
Temperature Range	(see specifications for drawers or enclosures)			
Humidity (Noncondensing)	(see specifications for drawers or enclosures)			
Wet Bulb Requirements	(see specifications for drawers or enclosures)			
Noise Emissions	(see specifications for drawers or enclosures)			
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	See physical planning diagrams for rack.			
1. Configuration-dependent, base weight plus weight of drawers				

Figure 117 on page 733 shows the physical planning diagram for the 7017 Models S70 and S70 Advanced Rack for AC systems.




Figure 118. 7017 Models S70 and S7A Rack Physical Planning Diagram

E.14 7014 Model S00 Rack Site and Hardware Planning Information

Table 231 summarizes the site and hardware planning information for the 7014 Model S00 Rack.

Table 231. 7014 Model S00 Rack Site and Hardware Planning Information

Dimensions				
Height	1577 mm		62.0 in	
Width	650 mm		25.5 in	
Depth	1019 mm		40.1 in	
Weight¹ (Base Rack)	159 kg		349 lbs	
Electrical	(see specifications for drawers or enclosures)			
Temperature Range	(see specifications for drawers or enclosures)			
Humidity (Noncondensing)	(see specifications for drawers or enclosures)			
Wet Bulb Requirements	(see specifications for drawers or enclosures)			
Noise Emissions	(see specifications for drawers or enclosures)			
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	1650 mm (65 in)	915 mm (36 in)	915 mm (36 in)	915 mm (36 in)
1. Configuration-dependent, base weight plus weight of drawers				

Figure 119 shows the physical planning diagram for the 7014 Model S00 Rack.




Figure 119. 7014 Model S00 Rack Physical Planning Diagram

E.15 7015 Model R00 Rack Site and Hardware Planning Information

Table 232 summarizes the site and hardware planning information for the 7015 Model R00 Rack.

Table 232. 7015 Model R00 Rack Site and Hardware Planning Information

Dimensions				
Height	1578 mm		62.0 in	
Width	650 mm		25.5 in	
Depth with standard door	921 mm		36.0 in	
Depth	1060 mm		41.8 in	
Weight ¹ (Base Rack)	130 kg		286 lbs	
Electrical	(see specifications for drawers or enclosures)			
Temperature Range	(see specifications for drawers or enclosures)			
Humidity (Noncondensing)	(see specifications for drawers or enclosures)			
Wet Bulb Requirements	(see specifications for drawers or enclosures)			
Noise Emissions	(see specifications for drawers or enclosures)			
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	1650 mm (65 in)	760 mm (30 in)	915 mm (36 in)	915 mm (36 in)
1. Configuration-dependent, base weight plus weight of drawers				

Figure 120 shows the physical planning diagram for the 7015 Model S00 Rack.




Figure 120. 7015 Model R00 Physical Planning Diagram

E.16 Model T00 and T42 Rack Site and Hardware Planning Information

Table 233 summarizes the site and hardware planning information for the Model T00 Rack, while Table 234 on page 742 refers to the Model T42 rack.

Table 233. Model T00 Rack Site and Hardware Planning Information

Dimensions				
Height	1804 mm			71.0 in
Capacity		36 EIA Units		
With PDP - DC Only	1926 mm			75.8 in
Width without side panels	623 mm			24.5 in
With side panels	644 mm			25.4 in
Depth with rear door	1042 mm			41.0 in
With both doors	1098 mm			43.3 in
Weight (Base Rack)	244 kg			535 lbs
Weight ¹ (Full Rack)	816 kg			1795 lbs
Electrical ²	(see specifications for drawers or enclosures)			
DC Rack				
Power source loading max ³		2.0 kVA		
AC Rack				
Power source loading per PDB ⁴		4.8 kVA		
Voltage range		200 to 240 V AC		
Frequency		50 or 60 Hz		
Temperature Range	(see specifications for drawers or enclosures)			
Humidity (Noncondensing)	(see specifications for drawers or enclosures)			
Wet Bulb Requirements	(see specifications for drawers or enclosures)			
Noise Emissions	(see specifications for drawers or enclosures)			
Clearances	Front	Back	Left	Right
Install/Air Flow	Maintenance of a proper service clearance should allow proper air flow.			
Service	1650 mm (65 in)	760 mm (30 in)	915 mm (36 in)	915 mm (36 in)
<p>1. Configuration-dependent, base weight plus weight of drawers</p> <p>2. The total rack power should be derived from the sum of the power used by the drawers in the rack</p> <p>3. The power distribution panel (PDP) on the DC powered rack can hold up to eighteen (nine per source) 48 volt 20 to 50 amp circuit breakers (configuration dependent).</p> <p>4. Each AC power distribution bus (PDB) can supply 4.8 kVA. A rack can have up to four PDBs as required by the drawers mounted in the rack.</p>				

Table 234. Model T42 Rack Site and Hardware Planning Information

Dimensions			
Height	2015 mm		79.3 in
Capacity		42 EIA Units	
With PDP - DC Only	N/A		N/A
Width without side panels	623 mm		24.5 in
With side panels	644 mm		25.4 in
Depth with rear door	1042 mm		41.0 in
With both doors	1098 mm		43.3 in
Weight (Base Rack)	261 kg		575 lbs
Weight [†] (Full Rack)	930 kg		2045 lbs

E.17 Noise Emission Notes

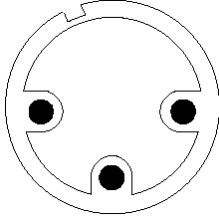

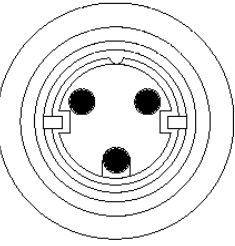
The following are noise emission notes that were referred to in the tables in this appendix:

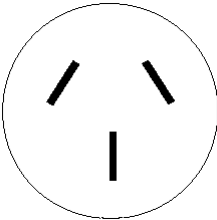
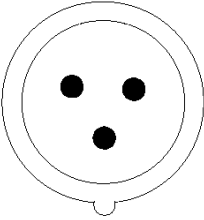
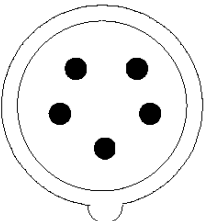
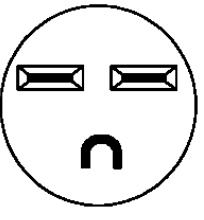
1. $L_{WA,d}$ is the declared sound power emission level for a production series of machines.
2. $L_{pA,m}$ is the mean value of the sound pressure emission levels at the operator position (if any) for a production series of machines.
3. $\langle L_{pA} \rangle_m$ is the mean value of the space-averaged sound pressure emission levels at the one-meter positions for a production series of machines.
4. N/A = Not Applicable (no operator position).
5. All measurements are made in accordance with ISO DIS 779 and reported in conformance with ISO DIS 7574/4.

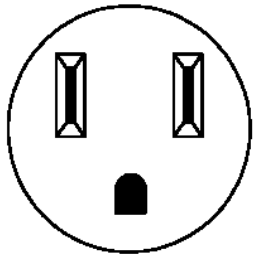
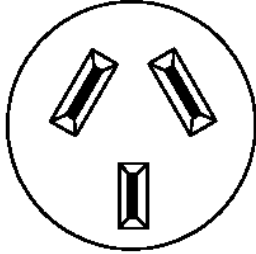
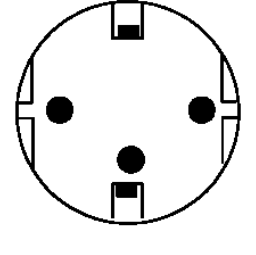
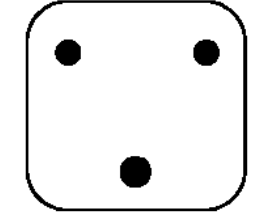
Appendix F. Power Cords

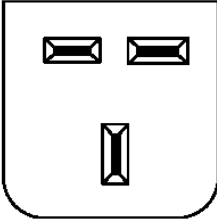


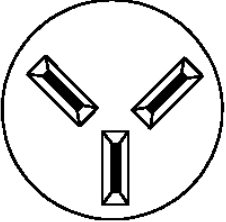
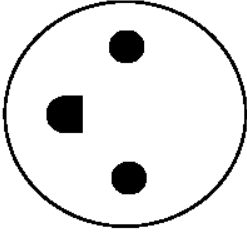
Table 235 provides a list of common power cords, a picture of their connectors (not to scale), and country information.

Table 235. Power Cord Connector Spotters Reference

Picture	Part Number	Description
	11F0106	For System Rack or Input/Output Rack, Wilco WP, type PDL, 250 V, single phase, 14ft., 30A Australia
	11F0107	For System Rack or Input/Output Rack, Wilco WP, type PDL, 250 V, single phase, 14ft., 30A, right angle New Zealand
	11F0113	For System Rack or Input/Output Rack, U.S. standard, type 12 plug, L6-30P, 250V, single phase, 14 ft., 30A Canada, U.S.A.
	11F0114	For System Rack or Input/Output Rack, U.S. Chicago, type 12 plug, L6-30P, 250V, single phase, 6 ft., 30A Chicago, Illinois, U.S.A.
	11F0115	For System Rack or Input/Output Rack, AFE, type 12 plug, L6-30P, 250V, single phase, 14 ft., 30A Argentina, Bahamas, Bangladesh, Barbados, Bermuda, Bolivia, Brunei, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Honduras, Hong Kong, India, Indonesia, Jamaica, Japan, Maco, Malaysia, Mexico, Myanmar, Netherlands Antilles, Panama, Peoples Republic of China, Peru, Philippines, Singapore, Sri Lanka, Suriname, Uruguay, Venezuela
	46F4593	For System Rack or Input/Output Rack, U.S. Chicago, WP, type 40 plug, R&S 3750, 250V, single phase, 6 ft., 30A Chicago, Illinois U.S.
	46F4594	For System Rack or Input/Output Rack, U.S., WP, type 40 plug, R&S 3750, 250V, single phase, 14 ft., 30A U.S.A.

Picture	Part Number	Description
	87G6067	For System Rack or Input/Output Rack, type KP, right angle, 250V, single phase, 14 ft., 30A Korea
	21H7693	For System Rack or Input/Output Rack, IEC 309, type 46 (2P+G), 250 V, single phase, 14ft., 32A New Zealand Belgium, Bahrain, Egypt, France, Greece, Iceland, Iraq, Ireland, Italy, Jordan, Kuwait, Lebanon, Malawi, Norway, Oman, Portugal, Qatar, Saudi Arabia, Spain, Turkey, U.K., United Arab Emirates
	88G4763	For Input/Output Rack, IEC 309, type 46 (3P+N+G), 250 V, two of three phase, 14ft., 32A Austria, Czech Republic, Denmark, Estonia, Finland, Israel, Latvia, Lithuania, Netherlands, Norway, Pakistan, Portugal, Russia, South Africa, Sweden, Turkey
	88G4764	For Input/Output Rack, IEC 309, type 46 (3P+N+G), 250 V, three phase, 14ft., 16A Switzerland
	1838574	Bahamas, Barbados, Bolivia, Brazil, Canada, Costa Rica, Dominican Republic, El Salvador, Ecuador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Netherlands Antilles, Panama, Peru, Philippines, Taiwan, Thailand, Trinidad Tobago, U.S.A., Venezuela

Picture	Part Number	Description
	6952301	Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Korea (South), Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Puerto Rico, Saudi Arabia, Suriname, Trinidad, Taiwan, U.S.A. (except Chicago), Venezuela
	13F9940	Argentina, Australia, New Zealand
	13F9979	Abu Dhabi, Austria, Belgium, Bulgaria, Botswana, Egypt, Finland, France, Germany, Greece, Iceland, Indonesia, Korea (South), Lebanon, Luxembourg, Macau, Netherlands, Norway, Portugal, Saudi Arabia, Spain, Sudan, Sweden, Turkey, Yugoslavia
	14F0015	Bangladesh, Burma, Pakistan, South Africa, Sri Lanka

Picture	Part Number	Description
	14F0033	Bahrain, Bermuda, Brunei, Channel Islands, Cyprus, Ghana, Hong Kong, India, Iraq, Ireland, Jordan, Kenya, Kuwait, Malawi, Malaysia, Nigeria, Oman, People's Republic of China, Qatar, Sierra Leone, Singapore, Tanzania, Uganda, United Arab Emirates (Dubai), United Kingdom, Zambia
	14F0051	Liechtenstein, Switzerland
	14F0069	Chile, Ethiopia, Italy
	14F0087	Israel
	13F9997	Denmark

Appendix G. Supported System Options

Table 236, Table 237 on page 748, and Table 238 on page 749 summarize the external devices that are supported by current RS/6000 models.

Table 236. External Devices for RS/6000 Models 140, 150, 170, 260, 270, and B50

	140	150	170	260	270	B50
Disk drives and subsystems						
7203-001 Portable Disk	X	X	X	X	X	
7204-409/419 External Disk	X	X	X	X	X	
7133 T40/D40 SSA Disk Subsystems	X	X	X	X	X	X
2102-D00 Expandable Storage Unit	-	-	-	-	-	
2102-F10 Fibre Channel RAID Storage Server	-	-	-	-	-	
2104-DL1/TL1 Expandable Storage Plus	X	X	X	X	X	X
2105-E10/E20 Enterprise Storage Server	-	-	-	-	-	
Optical drives and libraries						
3995-Cxx Optical Library	X	X	X	X	X	
7210-020 CD-ROM drive	X	X	X	X	X	
Tape drives and libraries						
7205-311 DLT Tape	X	X	X	X	X	X
7206-005 4 mm Tape	X	X	X	X	X	X
7206-110 4 mm Tape	X	X	X	X	X	X
7207-122 4 GB 1/4-inch Tape Cartridge	X	X	X	X	X	X
7207-315 4 GB 1/4-inch Tape Cartridge						X
7208-341 8 mm Tape	X	X	X	X	X	X
3490E-Fxx 1/2-inch Tape Subsystem (18/36-track)	X	X	X	X	X	X
3494-L12/L10/D12/D10 Magstar Tape Library	X	X	X	X	X	
3494-B18 Magstar Virtual Tape Server	-	-	-	-	-	
3570-Bxx Magstar MP Tape Subsystem	-	-	-	-	-	-
3570-Cxx Magstar MP Tape Subsystem	X	X	X	X	X	X
3575 Magstar MP Tape Library Dataserver	X	X	X	X	X	X
3580 Ultrascalable Tape Library	X	X	X	X	X	X
3590-E11/B11/C12 Magstar Tape	X	X	X	X	X	X
7331-305 8 mm Tape Library	X	X	X	X	X	X
7332-110 4 mm DDS-3 Tape Autoloader	X	X	X	X	X	
7332-220 4 mm DDS-3 Tape Autoloader						X
7337-305/306 DLT Tape Library	X	X	X	X	X	X
12 GB/24 GB 4 mm Tape (internal)	X	X	X	X	X	
5 GB/10 GB 8 mm Tape (internal)	-	-	-	-	-	
20/40 GB 8 mm Tape (internal)	-	-	X	X	X	
Communications subsystems						

	140	150	170	260	270	B50
Disk drives and subsystems						
8361-100/200 IBM Network Station	X	X	X	X	X	X

x = available or standard feature

Table 237. External Devices for RS/6000 Models F50, F80, H50, H70, and M80

	F50	F80	H70	H80	M80
Disk drives and subsystems					
7203-001 Portable Disk	X	X	X	X	X
7204-409/419 External Disk	X	X	X	X	X
7133 T40/D40 SSA Disk Subsystems	X	X	X	X	X
2102-D00 Expandable Storage Unit	X		X	X	X
2103-H07 Fibre Channel Storage Hub				X	X
2102-F10 Fibre Channel RAID Storage Server	X		X	X	X
2104-DL1/TL1 Expandable Storage Plus	X	X	X	X	X
2105-E10/E20 Enterprise Storage Server	X	X	X	X	X
Optical drives and libraries					
3995-Cxx Optical Library	X	X	X	X	X
7210-020 CD-ROM drive	X	-	X	-	-
Tape drives and libraries					
7205-311 DLT Tape	X	X	X	X	X
7206-005 4 mm Tape	X	X	X	X	X
7206-110 4 mm Tape	X	X	X	X	X
7207-122 4 GB 1/4-inch Tape Cartridge	X	X	X	X	X
7208-341 8 mm Tape	X	X	X	X	X
3490E-Fxx 1/2-inch Tape Subsystem (18/36-track)	X	X	X	X	X
3494-L12/L10/D12/D10 Magstar Tape Library	X	X	X	X	X
3494-B18 Magstar Virtual Tape Server	X	-	X	-	-
3570-Cxx Magstar MP Tape Subsystem	X	X	X	X	X
3575 Magstar MP Tape Library Dataserver	X	X	X	X	X
3580 Ultrascalable Tape Library	X	X	X	X	X
3590-E11/B11/C12 Magstar Tape	X	X	X	X	X
7331-305 8 mm Tape Library	X	X	X	X	X
7332-110 4 mm DDS-3 Tape Autoloader	X	X	-	X	X
7337-305/306 DLT Tape Library	X	X	X	X	X
12 GB/24 GB 4 mm Tape (internal)	X	X	X	X	X
5 GB/10 GB 8 mm Tape (internal)	X	-	X	-	-
20/40 GB 8 mm Tape (internal)	X	X	X	X	X
Communications subsystems					
8361-100/200 IBM Network Station	X	X	X	X	X

x = available or standard feature

Table 238. External Devices for RS/6000 Models S7A, S80, and SP

	S7A	S80	SP
Disk drives and subsystems			
7203-001 Portable Disk	X	X	-
7204-409/419 External Disk	X	X	-
7133 T40/D40 SSA Disk Subsystems	X	X	X
2102-D00 Expandable Storage Unit	X ¹	X ¹	-
2102-F10 Fibre Channel RAID Storage Server	X	X	-
2104-DL1/TL1 Expandable Storage Plus	X	X	X
2105-E10/E20 Enterprise Storage Server	X	X	X
Optical drives and libraries			
3995-Cxx Optical Library	X	X	X
7210-020 CD-ROM drive	X	X	X
Tape drives and libraries			
7205-311 DLT Tape	X	X	-
7206-005 4 mm Tape	X	X	-
7206-110 4 mm Tape	X	X	X
7207-122 4 GB 1/4-inch Tape Cartridge	X	X	X
7208-341 8 mm Tape	X	X	-
3490E-Fxx 1/2-inch Tape Subsystem (18/36-track)	X	X	X
3494-L12/L10/D12/D10 Magstar Tape Library	X	X	X
3494-B18 Magstar Virtual Tape Server	X	X	-
3570-Cxx Magstar MP Tape Subsystem	X	X	X
3575 Magstar MP Tape Library Dataserver	X	X	X
3590-E11/B11/C12 Magstar Tape	X	X	X
7331-305 8 mm Tape Library	X	X	X
7332-110 4 mm DDS-3 Tape Autoloader	X	X	-
7337-305/306 DLT Tape Library	X	X	X
12 GB/24 GB 4 mm Tape (internal)	X	X	-
5 GB/10 GB 8 mm Tape (internal)	X	X	X
20/40 GB 8 mm Tape (internal)	X	X	X
Communications subsystems			
8361-100/200 IBM Network Station	X	X	-

x = available or standard feature

Notes:

1. Not supported when the S70 Advanced/S80 is attached to the SP.

Appendix H. IBM InfoWindow II 3153

The InfoWindow 3153 is a popular choice as a system console for RS/6000s that do not come with a graphics accelerator. This appendix discusses the attributes of the IBM InfoWindow II 3153.

- 14-inch, 11.0-inch viewable image size, monochrome display (green, amber) with select-able overscan border that provides bezel-to-bezel video, offering screen formats of 24, 25, 28, 42, 43 or 50 data rows in any combination of 80 or 132 columns
- Eleven ASCII/ANSI/PC emulations:
 - IBM 3151
 - WYSE 50/50+, 60
 - DEC VT320, VT220, VT100
 - TeleVideo 925
 - ADDS Viewpoint
 - SCO Console
 - AT386
 - PC Term
- Two serial communications ports and one parallel printer port
- Main serial port supporting a speed of 134.4 Kbps - with a choice of RS232C (3153 Model BG3,BA3) or RS422A (3153 Model CG3,CA3) communications interface
- Auxiliary serial port supporting a speed of 38.4 Kbps using RS232C communications interface
- Simultaneous dual session by dual port
- Split Screen
- Up to ten pages of memory
- 28 programmable keys
- Copy-Paste
- Programmable Banner
- Desktop accessories: Calendar, Calculator, Clock, ASCII Chart, and Diagnostic Menu
- Power Management
- 30 Keyboard Languages
- 6 Setup Menu Languages

- 19 7-bit National Replacement Character Sets and 18 8-bit Multinational Code Pages
- PC-style or ASCII-style keyboard
- Compliant with ISO 9241 Part 3, MPR-II (magnetic and electric field), and EPA Energy Star
- Tilt and Swivel base, small footprint
- Limited three-year warranty

Appendix I. Customer Installation Matrix and Processor Groups

Table 239 shows which models are intended to initially be set up by the customer and which are intended to be set up by the IBM Customer Engineer/Customer Service Representative (CE/CSR). It also lists which features are intended to be installed by the customer and which features are to be installed by a CE/CSR as part of a Miscellaneous Equipment Specification (MES).

Table 239. RS/6000 Models and Feature Codes Set Up Information

Machine Type	Processor Group	Model	Initial System Customer Set-Up	MES Features/Options	
				CE Install	Customer Install
7006	D5	All	Yes	All features	None
7007		All	Yes	All features	None
7008		All	Yes	All features	None
7009		All	Yes	All features	None
7010		All	Yes	All features	None
7011	D5	All	Yes	All features	None
7012	E5	All	Yes	All features	None
7013	F5	All	No	All features	None
7015	G5	All	No	All features	None
7017	G5	All	No	All features	None
7024	D5	All	Yes	6309	All other features
7025	E5	F30	Yes	6309	All other features
7025	E5	F40	Yes	2856, 6309, 6549	All other features
7025	D5	F50	Yes	6309	All other features
7025	E5	F80	Yes	6309	All other features
7026	E5	All -M80	No	All other features	3083
7026	F5	M80	No	All features	None

Machine Type	Processor Group	Model	Initial System Customer Set-Up	MES Features/Options	
7027		All	No	All other features	2616, 3080, 3083, 3084, 3090, 6142, 6147, 3133, 3134, 3137, 3138, 6153, 6294, 6295
7043	D5	All (except 270)	Yes	6309	All other features
7043	E5	270	No	6309	All other features
7044	D5	170	Yes	6309	All other features
7044	E5	270	Yes	6309	All other features
7046	D5	B50	Yes	6309	All other features
7236		All	No	All features	None
7248	D5	All	Yes	2856	All other features
7317	E5	All	No	All features	None
7318		All	No	All features	None
7319		All	No	All features	None
9076	P5	All	No	All features	None

Appendix J. The RS/6000 64-Bit Solution

IBM RS/6000 brings 64-bit technology to the market with the introduction of the RS/6000 Enterprise Server Model S70 and AIX Version 4.3. Together, they provide the key elements of the 64-bit computing environment, both hardware and software. This section describes 64-bit technology and its value as well as the strategy and implementation for the RS/6000 64-bit solution. The essence of the RS/6000 64-bit computing strategy can be summed up in three themes, which will be elaborated upon throughout this section:

1. 64-bit is a new dimension of scalability. For the RS/6000, it is an additional dimension, complementing the established scalability of the existing 32-bit product set. 64-bit technology is the enabler for scaling enterprise SMP servers to higher capacity, making high-end system performance one of the primary customer benefits of 64-bit computing. This enabling may be summarized as: 64-bit hardware expands addressability, addressability creates the potential for higher capacity, and a balanced 64-bit system design will realize the potential by delivering higher system performance.
2. 64-bit computing is complementary to 32-bit computing. Customers want the benefits of 64-bit technology available to them, but know their 32-bit systems and 32-bit applications will be important investments for a very, very long time. A strategic focus on the complementary nature of the new technology forms the thread connecting many RS/6000 and AIX Version 4.3 64-bit features. It is also tightly connected to the third theme.
3. The transition from 32-bit computing to a future in which 64-bit and 32-bit computing coexist can be, and will be, for RS/6000 customers, a very smooth evolution.

These themes were also the design principles behind the product implementation of AIX Version 4.3 as a 64-bit operating environment. While the focus of this section is RS/6000 64-bit computing, it is of primary importance to note that AIX Version 4.3 also introduces significant functional and scalability enhancements that benefit all RS/6000 customers. A single AIX product supporting both 64-bit computing as well as broad general improvements is a prime example of the RS/6000 evolutionary vision at work.

J.1 What is 64-Bit Computing?

There is no formal definition of 64-bit computing. Its distinguishing technical features, as contrasted with 32-bit computing, are:

- Large file support
- Very large memory support

- Large application virtual address spaces
- 64-bit integer computation, using hardware with 64-bit general purpose registers.

An emphasis on the unique characteristics of 64-bit computing does not imply it is mutually exclusive with 32-bit computing. The picture of the future drawn in this section would more accurately be called a future of 64-bit and 32-bit coexistence. 64-bit Computing Characteristics

J.1.1 Large File Support

The ability to address data in files larger than 2 GB requires that a program be able to specify file offsets larger than a 32-bit number. This capability is generally considered to be 64-bit computing function, even though it does not require 64-bit hardware support. AIX Version 4.2 provided this capability for 32-bit programs, and AIX Version 4.3 provides it for 64-bit programs as well. Since it does not depend on 64-bit hardware, this function can be used on any RS/6000 system running the appropriate release of AIX.

There is, however, synergy between large file support and 64-bit hardware capabilities, in that a 64-bit program can have much larger portions of 64-bit files in its address space, as well as in system memory, at one time, than a 32-bit system could provide.

J.1.1.1 Large Physical Memory Support

Sufficient system memory is crucial for sustaining overall system performance. As a system's processor capacity grows with the speed and number of processors in the system, so does the requirement for system memory. Memory is a key element of having balanced resources, a requirement that applies whether the workload consists of 32-bit or 64-bit applications. For many customer environments, system memory capacity beyond 4 GB will be needed for optimum performance on an 8-way Model S70, with even more needed on a 12-way configuration. This is true even when the entire workload consists of growth in the throughput of 32-bit programs, which is one of the scalability scenarios the Model S70 is designed to support.

The new dimension of scalability introduced by 64-bit technology is the opportunity for some programs to keep very large amounts of data in memory, both resident in physical memory and accessible in their 64-bit virtual memory address space. While exploiting this new capability can significantly improve performance for some applications, there are relatively few types of applications that have evolved to make use of such techniques today. Those

that have exploited this capability, however, most often make use of very large memory for just one application program.

J.1.1.2 Large 64-bit Application Virtual Address Spaces

In 32-bit systems, an individual program, or process, may typically have between 2 GB and 4 GB of virtual address space for its own use to contain instructions and data. With 64-bit computing, applications may run in a 64-bit address space, where an individual program's addressability becomes measured in terabytes (TB).

Types of applications that currently understand how to exploit this opportunity include:

- Some database management programs use a large address space for scalability to maintain very large data buffers in memory, thus reducing the amount of I/O they need to perform. Using a large address space, they can supply data to client applications at the pace needed to sustain the high transaction rate potential afforded by many of the new processors in the industry.
- In certain cases, database management programs or customer applications may benefit from keeping an entire database or large file immediately accessible in memory. Read-only data lends itself most readily to this scenario. Significant improvements in response time or transaction rates are possible.
- Certain types of applications are able to directly attack larger problems by organizing larger arrays of data to be computed upon. Computer simulation of a physical phenomenon, such as aircraft flight or a nuclear reaction, are frequently cited examples.

J.1.1.3 64-bit Integer Computation

Native 64-bit integer computation is provided by 64-bit hardware and is utilized by programs computing on 64-bit data types. While there are some applications that need to do computation on integer numbers larger than 2^{32} , the key benefit of this capability is in performing arithmetic operations on pointers in 64-bit programs. Floating point computation already includes 64-bit precision on all RS/6000 systems.

J.1.1.4 64-bit Standards

Two distinct topics are included in the term *64-bit standards*. One is the latest step in the evolution of open systems standards, the Single Unix Specification Version 2, also referred to as Unix 98; the second is a broad industry agreement on data sizes in 64-bit programs, referred to as the LP64 data model.

The Single Unix Specification Version 2 takes account of 64-bit computing, without actually defining it, and does not specify any dependency on 64-bit hardware. The specification defines a programming environment for large files in a way that allows both 32-bit and 64-bit programs to be provided with this capability. The specification also cleans up a few APIs (application programming interfaces) that carried implications of 32-bit data types, therefore, allowing a 64-bit programming environment the same opportunity to be standards conformant as a 32-bit programming environment. The nature of the specification is such that a conforming 64-bit environment might exist on a system that supports only 64-bit programs or on a system that supports other environments (such as 32-bit binary compatibility) in addition to 64-bit programs.

J.1.1.5 The Need for 32-bit Compatibility and Interoperability

If one considers the benefits of 64-bit computing in the context of a customer with an established information technology investment based on existing 32-bit system implementations, it seems obvious that 64-bit computing will complement 32-bit computing in different ways for different customers, and the customers will move at a wide range of speeds into exploitation of the various elements of 64-bit computing. Large memory can benefit 32-bit application workloads; large file support can be, and has been, introduced for use by both 32-bit and 64-bit programs. The 64-bit virtual address space is of value to specific applications that can exploit it. Other applications are best left as 32-bit programs, but must be able to coexist and interoperate with the selective, but typically strategic, applications that are ported to 64-bit.

There is no single dimension to scalability. Some environments can grow with more 32-bit processes running on larger SMP systems. In many server environments, a critical server application must also scale with system throughput using larger files or bigger data buffers large enough to demand a 64-bit address space. Customer investment protection assistance, combined with customer flexibility to implement 64-bit technology as their environment demands it, leads to the requirement that 64-bit systems should support 32-bit binary compatibility, and that 32-bit and 64-bit environments must coexist and cooperate and share resources as easily as 32-bit programs traditionally have in the past.

J.1.2 Compatibility Starts at the Hardware Level

The following sections discuss hardware compatibility.

J.1.2.1 Hardware Binary Compatibility

RS/6000 64-bit products are built upon the PowerPC processor architecture, which defines both 64-bit and 32-bit processors, the latter as a subset of the

former. In practical terms, this means a PowerPC 64-bit processor is a proper superset of a 32-bit processor, therefore, allowing a 64-bit processor to run 32-bit programs the same way the programs run on a 32-bit processor. Hardware 32-bit binary compatibility is native; it does not rely on high-overhead emulation techniques. Hardware support for 32-bit binary compatibility is the first step in achieving 32-bit application binary compatibility on 64-bit systems under AIX Version 4.3. This hardware compatibility is also used by AIX itself, as a building block in maintaining binary compatibility for existing device drivers and kernel extensions running on 64-bit systems.

J.1.2.2 64-bit Hardware Characteristics

With full hardware 32-bit binary compatibility as the baseline, the features that characterize a PowerPC processor as 64-bit include:

- 64-bit general registers
- 64-bit instructions for loading and storing 64-bit data operands and for performing 64-bit arithmetic and logical operations.
- Two execution modes: 32-bit and 64-bit. While 32-bit processors have implicitly only one mode of operation, 32-bit execution mode on a 64-bit processor causes instructions and addressing to behave the same as on a 32-bit processor. As a separate mode, 64-bit execution mode creates a true 64-bit environment with 64-bit addressing and instruction behavior.
- 64-bit physical memory addressing facilities additional supervisor instructions as needed to set up and control the execution mode. A key feature the PowerPC 64-bit architecture provides is execution mode on a per-process level, helping AIX to create, at the system level, a mixed environment of concurrent 32-bit and 64-bit processes.

J.1.3 The RS/6000 Enterprise Server Model S70

The first 64-bit RS/6000 hardware system is the Model S70, available with 4, 8, or 12 processors, up to 16 GB of memory, and up to 56 PCI I/O slots. The S70 supports all the attributes of 64-bit computing and all the 32-bit binary compatibility features described in previous sections. While the 64-bit processors and very large memory provide obvious contrast to 32-bit systems, the Model S70 compatibly extends the scalability of the RS/6000 product line in other areas with features such as hot-swappable disks, online AIX diagnostics, a service processor that can detect system trouble and report it to the customer or to service professionals, and PCI I/O adapters and devices that include both SSA and SCSI disks.

J.1.4 AIX Version 4.3 - Single Product Strategy

The role of AIX in the RS/6000 64-bit strategy is completely grounded in the premise of an evolutionary transition to 64-bit computing. As AIX 4.3 introduces 64-bit computing, it remains one product that supports the full spectrum of current RS/6000 systems, both 32-bit and 64-bit. A broad range of functional enhancements not related to 64-bit are also available in AIX Version 4.3 on both 32-bit and 64-bit systems. On RS/6000 64-bit systems, these improvements are available to both 32-bit and 64-bit programs with AIX 4.3.

While 64-bit technology might appear to offer the opportunity for a multiple product strategy and the potential for divergence between 64-bit and 32-bit versions of software, AIX Version 4.3 builds directly on top of the established strengths of prior releases by integrating 64-bit support into a single product strategy, as AIX has done in the past through previous hardware architecture evolution. In particular, the introduction of 64-bit computing does not change AIX's commitment to:

- Release-to-release binary compatibility for customers with 32-bit systems. This applies, as it has in the past, both to applications and kernel extension or device driver programs (with caveats to cover specific, non-well-performing situations).
- Existing and emerging industry standards. AIX Version 4.3 is designed to provide the function specified in the Open Group Single Unix Specification Version 2, known informally as Unix 98. Note that open systems standards have not moved to 64-bit by leaving 32-bit behind; they have been expanded and generalized as needed to allow both 32-bit and 64-bit systems to comply. AIX Version 4.3 emulates this model by providing the same level of function on both 32-bit and 64-bit RS/6000 systems, and on 64-bit systems, the same AIX functions exist for both the 32-bit and 64-bit programming environments.

J.1.4.1 AIX Version 4.3 64-bit Implementation Strategy

The implementation philosophy in AIX for supporting 64-bit systems is to provide 32-bit binary compatibility, therefore, utilizing the hardware capability for 32-bit execution mode plus AIX release-to-release binary compatibility, providing coexistence and interoperability between 32-bit and 64-bit programs, and exploiting this infrastructure within AIX 4.3 itself.

The approach in AIX is no different from that recommended to, and likely to be adopted by, most third parties and customers as well. Most elements of AIX are not related in any way to supporting 64-bit hardware or providing 64-bit features, and many are small programs with modest address space

requirements. These programs remain 32-bit programs, with the same binary running on both 32-bit hardware or 64-bit hardware. Even in the cases where such programs were changed in AIX Version 4.3 to incorporate improved function or performance, they continue to be 32-bit programs. The coexistence of 32-bit and 64-bit programs on 64-bit hardware works for AIX itself, just as it will work for customers and third parties.

J.1.4.2 32-bit Binary Compatibility on 64-bit Hardware

Binary compatibility for 32-bit programs on 64-bit systems means, quite simply, that 32-bit applications run on 64-bit hardware systems the same as they would under AIX 4.3 on 32-bit hardware, and that the same binary executes the same hardware instructions on both systems. There is only one application binary required. AIX 4.3 kernel services for 32-bit applications look the same on 64-bit systems as on 32-bit systems - there is no compatibility overhead layer. Much of the AIX Version 4.3 kernel exploits the Model S70's 32-bit hardware binary compatibility to run in 32-bit execution mode. This helps AIX 4.3 to provide a native interface into kernel services from 32-bit applications.

Kernel exploitation of hardware binary compatibility also helps AIX Version 4.3 provide binary compatibility for well-performing 32-bit device drivers and kernel extensions running on 64-bit systems with 32-bit application workloads. There are a few situations, described later in this section, where an existing device driver or kernel extension must be enhanced to handle direct interaction with a 64-bit application, but this can be done in ways that do not impact its compatibility for 32-bit applications.

Under the RS/6000 philosophy of not forcing unnecessary recompilation to 64-bit, the large majority of application programs are likely to remain 32-bit, while a relatively small number of key programs scale up to exploit the performance benefits of a 64-bit address space. 32-bit binary compatibility in AIX Version 4.3, on both 32-bit and 64-bit systems, is an important ingredient for building the smooth, evolutionary road to 64-bit computing.

J.1.4.3 64-bit and 32-bit Interoperability Requirements

32-bit binary compatibility on 64-bit systems is not sufficient, however. It is not enough to say that 32-bit programs will continue to run on the same system where 64-bit programs also run. Serious investment protection means being able to leave 32-bit programs alone, even when another program they cooperate with is changed to exploit a 64-bit address space. A simple example of interoperability is provided by the sophisticated interaction between processes in a transaction processing environment. Client processes provide end-user transaction function, calling on the database

management process(es) for reading or updating database objects. The database manager program may become a 64-bit application process to expand its buffer space into very large memory. The client programs that call the database manager should continue to run as 32-bit programs, with binary compatibility, and no requirement to recompile because the database manager changed. The operating system's role in this example is to provide the infrastructure by which this goal is easily achieved by the middleware applications.

The open systems standards do not define interoperability requirements between 32-bit and 64-bit programming environments. The extensive interoperability in AIX Version 4.3 is the result of viewing 32-bit and 64-bit programs as equals, differentiated ONLY by their address space sizes and data model sizes (ILP32 versus LP64), and making it a design requirement in AIX Version 4.3 that they otherwise cooperate and interoperate as well as homogeneous processes do.

The use of the term *two programming environments* in this section refers to two environments in which programs can run on RS/6000 64-bit systems with AIX Version 4.3. Table 242 shows these AIX programming environments.

Table 240. AIX Programming Environments

AIX Version 4.3 Two Programming Environments	
32-bit environment	64-bit environment
4 GB address space	1 million TB address space
32-bit execution mode	64-bit execution mode
Data model ILP32 (32-bit pointer)	Data model LP64 (64-bit pointer)
XCOFF binary executable	XCOFF64 binary executable
32-bit shared libraries	64-bit shared libraries

Descriptions of these attributes are included in later sections.

J.1.4.4 64-bit and 32-bit Interoperability Implementation

64-bit and 32-bit interoperability refers to the two programming environments that AIX Version 4.3 provides on 64-bit systems, where 64-bit and 32-bit processes coexist, interact, and share common resources. This interoperability is actually present at several levels:

- Process-to-process interaction, the most obvious case.

- The infrastructure of 32-bit and 64-bit coexistence, which primarily means the creation, management, and location of coexisting 32-bit and 64-bit executables and shared libraries.
- Kernel support for two types of processes treated as peers

J.1.4.5 Process to Process Interaction

32-bit and 64-bit processes interact in the same ways that a collection of 32-bit and 64-bit processes may interact:

- They can fully share system resources - Files, shared memory, and interprocess communication (IPC) resources - and can signal each other.
- They can exec() each other without concern or awareness of whether the process being created will be 32-bit or 64-bit.
- They can set process resource limits that will apply to a process of the opposite type. The extensions in this area, for 32-bit processes to set 64-bit limits, are modeled upon comparable extensions provided for large file support in a 32-bit environment.

J.1.4.6 The Infrastructure of 32-bit and 64-bit Coexistence

The general strategy in infrastructure areas is to extend 32-bit programs with 64-bit awareness and function. The programs themselves generally do not benefit from a large address space; keeping them 32-bit has the added benefit, in some cases, of extending parts of the 64-bit infrastructure onto 32-bit systems.

In the area of 64-bit application development and library management:

- The compilers, assembler, linker, and other tools that create or manage binary executables, are extended to handle the new XCOFF64 binary format in addition to retaining compatible operation for 32-bit XCOFF objects.
- The programs listed above remain 32-bit applications, which means they can run on any AIX Version 4.3 system - both 32-bit and 64-bit. 64-bit program development can be done on existing 32-bit systems running AIX 4.3, with only the testing and debugging of 64-bit programs requiring 64-bit hardware.
- AIX Version 4.3 header files have been enhanced to support either 32-bit or 64-bit program development from single source header files.
- The infrastructure of AIX Version 4.3 shared libraries has been enhanced to support the coexistence of both 32-bit and 64-bit shared objects. Library pathnames need not change when a program is recompiled from 32-bit to 64-bit.

- AIX Version 4.3 shared library locations have not changed from previous AIX Version 4 releases; both 32-bit and 64-bit shared libraries coexist at the same pathname location.
- Makefiles can remain generic, with compilation mode controlled by switches, without complicated second order effects, such as separate library pathnames.
- Maintaining single source programs (and makefiles) for dual compilation into both 32-bit and 64-bit executables is straightforward.
- For compatibility, tools that manage the contents of shared library files can restrict their visibility to 32-bit objects only. This view can be controlled by command switch or an environment variable. The default mode is to operate on only the 32-bit content of shared libraries, thereby providing compatibility for 32-bit environments (whether on 32-bit systems or on 64-bit systems running primarily 32-bit applications).

Other elements of the coexistence infrastructure include:

- 64-bit capable shells. The shell programs remain 32-bit applications. Their 64-bit capability is the ability to handle the large numbers associated with resource limit settings for 64-bit processes.
- Most of the low 4 GB of the 64-bit application virtual address space has been reserved for application use. Specifically, those areas of a 32-bit address space, which are available for program data, is also available in a 64-bit address space. With care, cooperating 32-bit and 64-bit processes may share data at the same address locations. Extreme care must be taken if the shared areas contain data of types 'long' or 'pointer' (the two data types that differ in size between 32-bit and 64-bit programs).
- An administrator may control the ability to run 64-bit programs on 64-bit systems. The default setting for this capability is on. This function may be useful when the intended usage of a 64-bit system is to provide increased capacity and performance for 32-bit workloads, and the administrator wants to restrict experimentation with 64-bit applications. This system setting does not restrict full exploitation of 64-bit physical memory by multiple 32-bit programs, nor does it restrict any 64-bit large file capability.

J.1.4.7 Kernel Support for Process Interoperability

On 64-bit systems, the AIX Version 4.3 kernel supports both 32-bit and 64-bit application virtual address spaces. An application address space is set up automatically based on the type of program being run. Numerous kernel components are involved, especially the virtual memory manager (vmm) and the system dispatcher, which sets the correct execution mode when dispatching an application thread.

The kernel also provides equal access to system services from both types of processes. Most of the knowledge about which type of process makes a request for system services ('system call') is encapsulated at the entry points into the kernel. Equal access to system services provides the process-to-process interoperability described earlier, with 32-bit and 64-bit processes having the same access to all system resources.

In most cases, the kernel infrastructure isolates device drivers and kernel extensions from direct interaction with processes. There are exceptions, and in these cases, the AIX Version 4.3 kernel adds infrastructure to support the dual process environments. Kernel extensions may dynamically add entry points into the kernel, callable from application processes. For unmodified existing kernel extensions, the default behavior is to create a kernel entry point for 32-bit processes. Kernel extensions that wish to add a kernel entry point from 64-bit processes can be modified to do this; the kernel extension would then support both 32-bit and 64-bit processes on 64-bit systems, and the same kernel extension binary should run on 32-bit systems with AIX Version 4.3 (where it would obviously never be called by a 64-bit process). The modification referred to need not be a code modification with recompile; in some cases, the modification of the kernel extension's configuration attributes would suffice.

A similar situation applies to device drivers when an application program uses the `ioctl()` interface, which passes application parameters directly to a device driver. If a parameter were a pointer or a structure containing a pointer or long, for example, the device driver would need to be changed to use additional care in interpreting the parameter. With binary compatibility, existing device drivers will run on 64-bit systems with AIX Version 4.3 and support `ioctl()` calls from 32-bit processes. An `ioctl()` call from a 64-bit process will be rejected with an error indication unless the device driver's attributes indicate it can accept `ioctl()` calls from 64-bit processes. To establish this attribute, some code modification, with recompile of the AIX 4.3 device driver, might be necessary. On the other hand, if a device driver's `ioctl()` parameters are always of the 'int' data type, which is of constant size in the two types of processes, no coding or recompile should be necessary; the change could be made with an update to the AIX 4.3 device driver's configuration attributes.

In all cases, the default behavior favors 32-bit binary compatibility. In the few situations where new 64-bit applications may not coexist transparently, the kernel provides extended infrastructure to kernel extensions and device drivers. Kernel extensions and device drivers that are part of AIX Version 4.3 support both 32-bit and 64-bit processes.

J.1.4.8 AIX Version 4.3 64-bit Application Support

Much of the definition of the 64-bit application environment has already been established by virtue of the interoperability and evolutionary migration strategies:

- A 64-bit program should have the same level of function available as does a 32-bit program.
- 64-bit programs interoperate with 32-bit programs.
- 64-bit and 32-bit processes are peers in the eyes of the AIX Version 4.3 kernel.

The distinctive elements of the 64-bit environment are those that relate directly to the virtual address space size, the 64-bit application data model (LP64), the new binary executable format (XCOFF64), and 64-bit execution mode on 64-bit hardware systems. Note that industry standards is not part of the differences. 64-bit programs do not have any separate or distinct standards applied to them, as discussed earlier.

J.1.4.9 64-bit Address Space

The new 64-bit application virtual address space is laid out in Table 241 on page 767. Immense regions are allocated for the different types of program and data areas that occupy an application address space. The range of addressability in the 64-bit address space is one million terabytes (TB), or one billion gigabytes (GB), or 2^{60} bytes. More than a third of this range has been reserved for application program text and data. Shared memory regions, shared libraries, dynamically loaded programs, and the program

stack are each allocated a subset of the address space up to 2^{56} bytes, or approximately 64,000 TB.

Table 241. AIX Virtual Address Space Layout

		Size	64-bit Address Space	32-bit Address Space
Top=1 Million TB X Above 4 GB		64,000 TB	Program Stack	N/A
		384,000 TB	Reserved	
		64,000 TB	Shared library text and data	
		64,000 TB	Privately loaded programs	
		64,000 TB	Shared objects (default)	
		384,000 TB	Program text and data	
Below 4 GB (Segment number)	15	256 MB	Reserved	Shared library data
	14	256 MB	Available for applications	Available for applications
	13	256 MB	Shared library information	Shared library text
	3-12	2.56 GB	Available for applications	Available for applications
	2	256 MB	Available for applications	Program data and stack
	1	256 MB	Available for applications	Program text
	0	256 MB	Kernel	Kernel

J.1.4.10 64-bit Data Model - LP64

The key attribute of 64-bit programs, from an application development perspective, is the difference in size of the C language data type's pointer and long. 64-bit pointers are an obvious necessity in a 64-bit address space. The

64-bit long type provides a basic integral arithmetic data type that corresponds to the computational size of hardware registers in 64-bit execution mode. Together they are referred to as the LP64 data model. The corresponding model for 32-bit programs is ILP32, that is, the types int, long, and pointer are each 32-bit quantities. In the LP64 model, only long and pointer size is different; other data types, including int, remain the same size as in the ILP32 model. The LP64 model is an industry convention that has been agreed to by system vendors, including IBM, Digital, SGI, HP, and Sun.

J.1.4.11 64-bit Executable Format - XCOFF64

AIX Version 4.3 support for 64-bit applications includes a new, additional executable file format, called XCOFF64. It is an extension of the XCOFF format used in AIX for 32-bit executable files. XCOFF64 extends executable file scalability in two ways:

- It provides for 64-bit addresses in the program code and data. Since the executable file's contents will run as a program in a 64-bit address space, XCOFF64 must provide for 64-bit address relocation.
- Offset and length values in the control areas of the file are enlarged to 64-bit quantities. The architecture of the XCOFF64 format provides for future growth in the size of 64-bit executables that is, in practical terms, virtually limitless.

The maximum size of a 64-bit executable file in AIX version 4.3 is 2 GB, which includes the code and initialized data of the programs in the file. This limit is imposed by current implementations of AIX Version 4.3 tools that manage XCOFF64 files; it is not an attribute of the XCOFF64 definition. Program uninitialized data, which does not consume space in the executable file, may be as large as approximately 1/2 million terabytes in a 64-bit program.

J.1.4.12 64-bit Execution Scope and 64-bit Shared Libraries

A 64-bit program executes as a 64-bit process, with a 64-bit address space, with the processor executing instructions in 64-bit execution mode. Inside a process, the execution environment is homogeneous, that is, all threads in the process execute in the same mode, and all programs called within the process execute in the same mode. The same is true for 32-bit processes. One of the obvious implications is that two versions of shared library programs must exist, one to run in 32-bit processes, and the other to run in 64-bit processes. (Note that homogeneity ends at the point where a process calls the operating system kernel; a single kernel supports both types of processes, and, by definition, cannot be homogeneous with both.)

Most key libraries shipped as part of AIX Version 4.3 have two binary versions. The 32-bit version will be used for linking with 32-bit programs, and the 64-bit version will be used for linking with 64-bit programs. This happens transparently, as discussed earlier in this section, since the two shared library entities coexist at the same pathname location.

Libraries and interfaces that are not provided in 64-bit format are primarily those that are obsolete with respect to compiling a new program or recompiling an old program to exploit the 64-bit address space. These libraries continue to be provided in 32-bit binary form to provide 32-bit binary compatibility for programs compiled and linked on older releases of AIX; programs being written for, or converted to, the 64-bit programming environment should be updated to current levels of standard APIs using AIX Version 4.3.

A small group of AIX Version 4.3 libraries remain in 32-bit format only but may be extended in the future to provide a 64-bit version. They are not obsolete, nor are they needed by the applications that are candidates for porting on the initial 64-bit system, the RS/6000 Enterprise Server Model S70. Examples include the graphical interfaces in the OpenGL and graPHIGS libraries.

J.1.4.13 AIX Version 4.3 64-bit Kernel Strategy

As AIX Version 4.3 introduces 64-bit computing, the primary issues for the AIX kernel are:

- Support for 64-bit computing functions.
- Equal support for both 32-bit and 64-bit processes.
- Compatibility for device drivers and other extensions that plug into AIX's highly modular and extensible kernel.
- Kernel scalability.
- Performance.

J.1.4.14 64-bit Computing Support

The AIX kernel implements the functions needed to provide a full 64-bit computing environment, supporting large memory, large files, 64-bit address spaces, and system services for 64-bit applications.

J.1.4.15 Support for 32-bit and 64-bit Process Coexistence

The kernel's role is to provide the same system services to both environments, treating the two environments as peers, creating the interoperability infrastructure by which 32-bit and 64-bit applications fully share system resources and interprocess communication facilities.

J.1.4.16 Compatibility

By enabling smooth customer migration to 64-bit systems, AIX Version 4.3 avoids requiring all new device drivers for the 64-bit environment. This is especially important in the RS/6000 context because:

- The complementary nature of 64-bit and 32-bit and the evolutionary RS/6000 strategy anticipate the likely situation that most applications running on 64-bit systems will continue to be 32-bit programs. This will change over time, but gradually. There is no good reason to require all new device drivers to run primarily preexisting 32-bit programs on new hardware, whether the new hardware is based on 64-bit technology or 32-bit technology.
- The AIX kernel has significant modularity and extensibility capabilities; so, there are a significant number of existing device drivers and other kernel extensions available from both IBM and other vendors as well as those packaged with the AIX operating system.

J.1.4.17 Kernel Scalability

The AIX kernel's existing scalability makes it analogous to a 32-bit application that does not require a 64-bit address space. The AIX Version 4.3 kernel can provide all the 64-bit computing function described above without encountering kernel address space constraint. This allows the AIX Version 4.3 kernel to adopt an evolutionary approach to internal 64-bit exploitation and retain 32-bit binary compatibility for device drivers and extensions. Kernel components make use of 64-bit instructions and 64-bit execution mode, as required, to provide full 64-bit computing function, insulated by the modularity provided by the AIX kernel program architecture.

J.1.4.18 Performance

Two important factors in AIX Version 4.3 kernel performance are program footprint and system call transition time. Kernel footprint is analogous to application program footprint. A 64-bit application creates a bigger footprint, in memory and in cache, than does a 32-bit program. If an application program needs 64-bit addressing, it is available and should be exploited. But if not, it is best to leave the program 32-bit. The same logic applies to the kernel. The kernel footprint is reduced as a result of reusing the existing, stable 32-bit kernel design.

With respect to system call transition time, there is one kernel supporting both 32-bit and 64-bit processes. In cases where pointers or complex data structures are passed to the kernel as a parameter, a small additional overhead may be incurred at the interface between the kernel and a process with dissimilar data types. In AIX Version 4.3, because of the evolutionary and

dual environment coexistence strategy, most processes on a system will be 32-bit processes. Most 64-bit processes will be 64-bit to consciously exploit the large memory and large address space features of 64-bit computing. In some cases, this will reduce the frequency of calls for system services due to a 64-bit process needing to perform less I/O or being able map large numbers of files directly into the address space. Overall, AIX Version 4.3 performance benefits are realized by optimizing the interface between the kernel and 32-bit processes in spite of the hardware being a 64-bit processor.

All of the above factors and considerations result in an implementation approach for the AIX Version 4.3 kernel that focuses on delivering 64-bit computing functions while reusing much of the existing kernel's 32-bit programming environment.

J.1.4.19 A Kernel Example: AIX Support for Very Large Memory

Support for large memory is one of the key elements of 64-bit computing. The approach to supporting large memory in the AIX Version 4.3 kernel serves as an example of how an evolutionary path has been adopted in implementing 64-bit function, therefore, complementing the existing AIX kernel base with 64-bit design.

Physical, or system, memory support is quite straightforward because it is all located within the AIX Version 4.3 kernel, and highly encapsulated in the component called the virtual memory manager (vmm). System memory is used, on any RS/6000 system, in the following ways:

- As the real memory pages that contain in-use pages of an application's virtual address space (called the application's *working set*).
- As the real memory pages containing kernel programs and data. Since large portions of the AIX kernel are pageable, real memory is needed for the kernel's in-use pages.
- As file system caching pages. The AIX Journaled File System (JFS) uses real memory to contain recently-used and read-ahead pages of data for files that applications have open.

On 64-bit systems, the AIX Version 4.3 vmm allocates the pages of very large memory for the same purposes. All memory pages, whether located at real addresses below or above 4 GB, can be used to sustain the working sets of 32-bit and 64-bit applications running concurrently. In some environments, large memory will support an expansion in the number of 32-bit applications. In other cases, a large portion of memory will contain the in-use pages for a single 64-bit program. Internally, the virtual memory manager in the AIX Version 4.3 kernel builds extended data structures to deal with the very large

number of pages typically present on a 64-bit system and deals with the 64-bit addresses of real memory. The scalability of the virtual memory manager's design is important, but not its execution mode within the kernel. The hardware translation facilities of the 64-bit processors are controlled by the virtual memory manager. I/O data transfers to 64-bit memory utilize new hardware DMA facilities, with kernel support insulated within the AIX DMA servers available to device drivers. In both cases, the hardware control facilities for 64-bit memory addressing can be managed using kernel code running in either 32-bit or 64-bit execution mode, based on the nature of the 64-bit hardware's 32-bit compatibility.

J.1.4.20 Summary

The IBM RS/6000 64-bit solution is introduced with the RS/6000 Enterprise Server Model S70 and AIX Version 4.3. An examination of the elements of 64-bit technology and the features of the RS/6000 64-bit products illustrates the value of a strategic view of 64-bit computing.

The RS/6000 strategy for 64-bit computing can be summarized with the same themes defined in the opening paragraphs: 64-bit is a new dimension of scalability. 64-bit computing is complementary to 32-bit computing. The expansion of RS/6000 computing breadth to include 64-bit computing will be a very smooth evolution.

Appendix K. Special Notices

This publication is intended to help IBM sales professionals, Business Partners, ISVs, and customers wishing to obtain a single source reference for IBM RS/6000 hardware and software offerings. The information in this publication is not intended as the specification of any programming interfaces that are provided by AIX software or by RS/6000 hardware. See the PUBLICATIONS section of the IBM Programming Announcement for RS/6000 and AIX LPPs 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, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate

them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	AIX/ESA
AIXwindows	FICON
Enterprise Storage Server	C Set++
GXT150L	OS/390
AS/400	DB2 Universal Database
ES/3090	DB2
ES/9000	DirectTalk
ESCON	HACMP/6000
OS/2	GXT150M
GXT1000	Portmaster
HAGEO	IBMLink
IBM	Netfinity Manager
LANStreamer	Micro Channel
LoadLeveler	Magstar
Netfinity	Network Station
POWER Gt3	POWER Gt3i

POWER Gt4
POWER Gt4i
POWER Gt4xi
POWER GTO
Power PC 601
Power PC 603
Power Series
POWER3 Architecture
POWERparallel
PowerPC Architecture
RS/6000
ServRAID
SP2
Ultrastar
VisualAge

POWER Gt4e
POWER Gt4x
POWER2 Architecture
POWER Gt1
Deep Blue
Power PC 604
PowerPC
RISC System/6000
Seascape
PowerPC Reference Platform
Scalable POWERparallel Systems
SP1
TXSeries
Versatile Storage Server

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere., The Power To Manage., Anything. Anywhere., TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks

owned by SET Secure Electronic Transaction LLC.

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

Appendix L. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

L.1 IBM Redbooks

For information on ordering these publications see “How to Get IBM Redbooks” on page 781.

- *RS/6000 43P 7043 Models 150 and 260 Handbook*, SG24-5144
- *RS/6000 Models E30, F40, F50, and H50 Handbook*, SG24-5143
- *RS/6000 Scientific and Technical Computing: POWER3 Introduction and Tuning Guide*, SG24-5155
- *Inside the RS/6000 SP*, SG24-5145
- *RS/6000 Graphics Handbook*, SG24-5130

L.2 IBM Redbooks Collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at ibm.com/redbooks for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
IBM System/390 Redbooks Collection	SK2T-2177
IBM Networking Redbooks Collection	SK2T-6022
IBM Transaction Processing and Data Management Redbooks Collection	SK2T-8038
IBM Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
IBM AS/400 Redbooks Collection	SK2T-2849
IBM Netfinity Hardware and Software Redbooks Collection	SK2T-8046
IBM RS/6000 Redbooks Collection	SK2T-8043
IBM Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

L.3 Other Resources

These publications are also relevant as further information sources:

- *IBM RS/6000 SP: Planning Volume 1, Hardware and Physical Environment*, GA22-7280
- *Exploring IBM RS/6000 Computers*, GA23-2674

- *IBM General Information Installation Manual - Physical Planning*, GC22-7072
- *IBM RS/6000 SP: Maintenance Information Volume 1*, GC23-3903
- *IBM RS/6000 SP: Maintenance Information Volume 2*, GC23-3904
- *Product Reference Guide, Open Systems Storage*, G325-3304
- *PCI Adapter Placement Reference*, SA38-0538
- *Site and Hardware Planning Information*, SA38-0508
- *RS/6000 Facts and Features*, G320-9878
- *PowerPC and POWER2 Technical Aspects*, SA23-2737

L.3.1 RS/6000 Business Partner Sales Kit CD

This handy CD is designed to provide 'one stop shopping' for a vast range of RS/6000 product and sales information.

The CD is an essential aid for RS/6000 sales and technical specialists within both IBM and the Business Partner communities. Providing fast, convenient and portable access, it complements other sources of information, such as the Internet and hardcopy.

The CD includes:

- Sales information
 - Presentations and brochures
 - Case studies and consultant reports
 - Competitive information
- Reference library
 - Detailed product descriptions
 - RS/6000 Redbooks
 - White papers
- Multimedia presentations

The CD is for IBM and authorized RS/6000 Business Partners only, and can be ordered through IBM Publications (SK2T-9772 in North America or SK2T-9771 in EMEA). The CD is updated twice per year, shortly after major announcements.

L.4 Referenced Web Sites

These Web sites are also relevant as further information sources:

AIX LPPs www.rs6000.ibm.com/software/Apps/LPPmap.html

Business intelligence systems

www.ibm.com/solutions/businessintelligence/solutions/index.htm

www.austin.ibm.com/software/Appfinder/decisionsupport.html

www.software.ibm.com/data/busn-intel

CATIA www.ibm.com/solutions/plm

Global Services www.rs6000.ibm.com/services

ISV solutions www.software.ibm.com/solutions/isv

Lotus Notes www.lotus.com/notes

LINPACK benchmark www.netlib.no.netlib/benchmark/performance.ps

NotesBench figures as.ideascp.com/cpweb

RS/6000 library www.rs6000.ibm.com/resource/hardware_docs

Solution segments www.rs6000.ibm.com/solutions/xxx
(where xxx is: bi, e-business, erp, oltp, or
supercomputing)

SPEC benchmark www.specbench.org

ESS Storage

<http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm#6>

How to Get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** ibm.com/redbooks

Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at <http://w3.ibm.com/> for redbook, residency, and workshop announcements.

IBM Redbooks Fax 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.

Abbreviations and Acronyms

ABI	Application Binary Interface	CATIA	Computer-Graphics Aided Three-Dimensional Interactive Application
AC	Alternating Current		
ADSM	ADSTAR Distributed Storage Manager	CD	Compact Disk
ADSTAR	Advanced Storage and Retrieval	CD-ROM	Compact Disk-Read Only Memory
AIX	Advanced Interactive Executive	CE	Customer Engineer
ANSI	American National Standards Institute	CEC	Central Electronics Complex
APAR	Authorized Program Analysis Report	CHRP	Common Hardware Reference Platform
API	Application Programming Interface	CLIO/S	Client Input/Output Sockets
ASCI	Accelerated Strategic Computing Initiative	CMOS	Complimentary Metal Oxide Semiconductor
ASCII	American National Standards Code for Information Interchange	COLD	Computer Output to Laser Disk
ATM	Asynchronous Transfer Mode	CPU	Central Processing Unit
BFF	Backup File Format	CRC	Cyclic Redundancy Check
BOS	Base Operating System	CSR	Customer Service Representative
BI	Business Intelligence	CSS	Communication Subsystems Support
BIST	Built-In Self-Test	CSU	Customer Set-Up
BLAS	Basic Linear Algebra Subprograms	CSU	Channel Service Unit
BOS	Base Operating System	CWS	Control Workstation
CAE	Computer-Aided Engineering	DAS	Dual Attach Station
CAD	Computer-Aided Design	DASD	Direct Access Storage Device (Disk)
CAM	Computer-Aided Manufacturing	DAT	Digital Audio Tape
		DC	Direct Current
		DDC	Display Data Channel
		DDS	Digital Data Storage
		DE	Dual-Ended

DFS	Distributed File System	FC-AL	Fibre Channel-Arbitrated Loop
DIMM	Dual In-Line Memory Module	FCP	Fibre Channel Protocol
DIP	Direct Insertion Probe	FDDI	Fiber Distributed Data Interface
DIVA	Digital Inquiry Voice Answer	FDX	Full Duplex
DLT	Digital Linear Tape	FRU	Field Replaceable Unit
DMA	Direct Memory Access	FTP	File Transfer Protocol
DOS	Disk Operating System	F/W	Fast and Wide
DRAM	Dynamic Random Access Memory	GPFS	General Parallel File System
DSU	Data Service Unit	GUI	Graphical User Interface
DW	Data Warehouse	HACMP	High Availability Cluster Multi Processing
EC	Engineering Change	HACWS	High Availability Control Workstation
ECC	Error Checking and Correction	HDX	Half Duplex
EEPROM	Electrically Erasable Programmable Read Only Memory	HIPPI	High Performance Parallel Interface
EIA	Electronics Industry Association	HIPS	High Performance Switch
EISA	Extended Industry Standard Architecture	HIPS LC-8	Low-Cost Eight-Port High Performance Switch
ELA	Error Log Analysis	HP	Hewlett-Packard
EMIF	ESCON Multiple Image Facility	HPF	High Performance FORTRAN
EPOW	Environmental and Power Warning	HPSSDL	High Performance Supercomputer Systems Development Laboratory
ESCON	Enterprise Systems Connection (Architecture, IBM System/390)	HP-UX	Hewlett-Packard UNIX
ESSL	Engineering and Scientific Subroutine Library	HTTP	Hypertext Transfer Protocol
ETML	Extract, Transformation, Movement and Loading	Hz	Hertz
F/C	Feature Code	IA	Intel Architecture

IBM	International Business Machines	LANE	Local Area Network Emulation
ID	Identification	LAPI	Low-Level Application Programming Interface
IDE	Integrated Device Electronics	LED	Light Emitting Diode
IDS	Intelligent Decision Server	LFT	Low Function Terminal
IEEE	Institute of Electrical and Electronics Engineers	LP	Linear Programming
I²C	Inter Integrated-Circuit Communications	LPP	Licensed Program Product
I/O	Input/Output	LVM	Logical Volume Manager
IP	Internetwork Protocol (OSI)	MAP	Maintenance Analysis Procedure
IPL	Initial Program Load	MAU	Multiple Access Unit
IrDA	Infrared Data Association (which sets standards for infrared support including protocols for data interchange)	Mbps	Megabits Per Second
		MBps	Megabytes Per Second
		MCA	Micro Channel Architecture
		MCAD	Mechanical Computer-Aided Design
IRQ	Interrupt Request	MES	Miscellaneous Equipment Specification
ISA	Industry Standard Architecture		
ISB	Intermediate Switch Board	MIP	Mixed-Integer Programming
ISDN	Integrated-Services Digital Network	MLR1	Multi-Channel Linear Recording 1
ISV	Independent Software Vendor	MMF	Multi-Mode Fibre
		MP	Multiprocessor
ITSO	International Technical Support Organization	MP	Multi-Purpose
JBOD	Just a Bunch of Disks	MPC-3	Multimedia PC-3
JFS	Journalled File System	MPI	Message Passing Interface
JTAG	Joint Test Action Group	MPP	Massively Parallel Processing
L1	Level 1		
L2	Level 2	MPS	Mathematical Programming System
LAN	Local Area Network		

MTU	Maximum Transmission Unit	PIOFS	Parallel Input Output File System
MVS	Multiple Virtual Storage (IBM System 370 and 390)	POE	Parallel Operating Environment
MX	Mezzanine Bus	POP	Power-On Password
NCP	Network Control Point	POSIX	Portable Operating Interface for Computing Environments
NFS	Network File System	POST	Power-On Self-test
NIM	Network Installation Manager	POWER	Performance Optimization with Enhanced Risc (Architecture)
NT-1	Network Terminator-1	PPP	Point-to-Point Protocol
NTP	Network Time Protocol	PREP	PowerPC Reference Platform
NUMA	Non-Uniform Memory Access	PSSP	Parallel System Support Program
NVRAM	Non-Volatile Random Access Memory	PTF	Program Temporary Fix
OCS	Online Customer Support	PTPE	Performance Toolbox Parallel Extensions
ODM	Object Data Manager	PTX	Performance Toolbox
OLAP	Online Analytical Processing	PV	Physical Volume
OS/390	Operating System/390	PVC	Permanent Virtual Circuit
OSL	Optimization Subroutine Library	QMF	Query Management Facility
OSLp	Parallel Optimization Subroutine Library	QP	Quadratic Programming
P2SC	Power2 Super Chip	RAM	Random Access Memory
PAP	Privileged Access Password	RAN	Remote Asynchronous Node
PBLAS	Parallel Basic Linear Algebra Subprograms	RAS	Reliability, Availability, and Serviceability
PCI	Peripheral Component Interconnect	RAID	Redundant Array of Independent Disks
PDU	Power Distribution Unit	RDBMS	Relational Database Management System
PE	Parallel Environment		
PEDB	Parallel Environment Debugging		
PID	Program Identification		

RIPL	Remote Initial Program Load	SMP	Symmetric Multiprocessing
ROLTP	Relative Online Transaction Processing	SOI	Silicon-on-Insulator
RPA	RS/6000 Platform Architecture	SP	Scalable POWERParallel
RVSD	Recoverable Virtual Shared Disk	SP	Service Processor
RTC	Real-Time Clock	SPEC	Standard Performance Evaluation Corp.
SAN	Storage Area Network	SPOT	Shared Product Object Tree
SAS	Single Attach Station	SPS	SP Switch
SAR	Solutions Assurance Review	SPS-8	Eight-Port SP Switch
ScaLAPACK	Scalable Linear Algebra Package	SRC	System Resource Controller
SCO	Santa Cruz Operations	SSC	System Support Controller
SCSI	Small Computer System Interface	SSA	Serial Storage Architecture
SDR	System Data Repository	STP	Shielded Twisted Pair
SDRAM	Synchronous Dynamic Random Access Memory	SUP	Software Update Protocol
SDLC	Synchronous Data Link Control	SVC	Switch Virtual Circuit
SE	Single-Ended	Tcl	Tool Command Language
SEPBU	Scalable Electrical Power Base Unit	TCP/IP	Transmission Control Protocol/Internet Protocol
SGI	Silicon Graphics Incorporated	TCQ	Tagged Command Queuing
SLIP	Serial Line Internet Protocol	TPC	Transaction Processing Council
SLR1	Single-Channel Linear Recording 1	UDB EEE	Universal Database and Enterprise Extended Edition
SMIT	System Management Interface Tool	UP	Uniprocessor
SMS	System Management Services	USB	Universal Serial Bus
		UTP	Unshielded Twisted Pair

UUCP	UNIX-to-UNIX Communication Protocol
VESA	Video Electronics Standards Association
VG	Volume Group
VM	Virtual Machine (IBM System 370 and 390)
VMM	Virtual Memory Manager
VPD	Vital Product Data
VSD	Virtual Shared Disk
VSM	Visual Systems Management
VSS	Versatile Storage Server
VT	Visualization Tool
WAN	Wide Area Network
WTE	Web Traffic Express
XTF	Extended Distance Feature

Index

Numerics

- 128-port asynchronous adapter
 - Remote Asynchronous Node (RAN) 333
- 140, Model 28
- 170, Model 40
- 2050 SP 332 MHz SMP Thin node 281
- 2051 SP 332 MHz SMP Wide node 270
- 2102-F10 IBM Fibre Channel RAID Storage Server 427
- 2104 LVD expandable storage plus 439
- 2493 SCSI-2 Fast/Wide RAID Adapter 349
- 2494 PCI 3-Channel Ultra SCSI RAID Adapter 350
- 24-way SMP computing 189
- 2639 Ultimedia Video Capture Adapter/S (PCI) 386
- 270, Model 53
- 2708 Eicon ISDN DIVA PRO 2.0 S/T Adapter 325
- 2741 FDDI SK-NET LP SAS Adapter 329
- 2742 FDDI SK-NET LP DAS Adapter 330
- 2743 FDDI SK-NET UP DAS Adapter 331
- 2751 S/390 ESCON Channel Adapter 326
- 2823 GXT2000P Graphics Accelerator 373
- 2825 GXT3000P Graphics Accelerator 369
- 2830 GXT130P Graphics Accelerator 381
- 2838 GXT120P Graphics Accelerator 384
- 2841 GXT300P Graphics Accelerator 376
- 2851 GXT250P Graphics Accelerator 379
- 2852 GXT255P Graphics Accelerator 379
- 2943 8-Port Asynchronous Adapter 210, 332
- 2944 128-Port Asynchronous Adapter 171, 203, 210, 333
- 2947 ARTIC960Hx 4-Port Selectable Adapter 335
- 2948 ARTIC960Hx 4-Port T1/E1 Adapter 336
- 2962 2-port Multiprotocol X.25 Adapter 338
- 2963 TURBOWAYS 155 PCI UTP ATM Adapter 340
- 2968 10/100 Ethernet TX PCI Adapter 342
- 2969 Gigabit Ethernet-SX PCI Adapter 344
- 2987 Ethernet AUI/RJ-45 Adapter 346
- 2988 TURBOWAYS 155 PCI MMF ATM Adapter 341
- 2998 TURBOWAYS 25 ATM Adapter 342
- 3125 serial-to-serial cable 171, 203, 210
- 3142 3-meter Remote I/O Cable, Drawer-to-Drawer 188
- 3143 6-meter Remote I/O Cable, Rack-to-Rack 188
- 3144 15-meter Remote I/O Cable, Rack-to-Rack 188
- 3153 InfoWindow 751
- 32-bit 509
- 32-bit PCI slots 189
- 332 MHz SMP Thin node 281
- 332 MHz SMP Wide node 270
- 3401 4.5 GB disk drive module 210
- 3466 Network Storage Manager 486
- 3490 tape units 195
- 3490E Magnetic Tape Subsystem 461
- 3494 tape library 466
- 3570 tape subsystem 471
- 3570 tape units 195
- 3570-B11 tape unit 195
- 3570-B12 tape unit 195
- 3575 tape library dataserver 473
- 3590 tape subsystem 463
- 3901 9.1 GB disk drive module 210
- 3995 optical library solutions 482
- 4007 SP High Performance Switch LC-8 Adapter 292
- 4008 SP Switch-8 Adapter 292
- 4010 High Performance Switch 292
- 4011 SP Switch 292
- 4022 SP Switch MX Adapter 292, 313
- 4023 SP Switch MX Adapter 292
- 43P Series 25
- 4951 4-Port 10/100 Base-TX Ethernet PCI Adapter 347
- 4959 4/16 Mbps token-ring adapter 348
- 5648-B64 performance toolbox parallel extensions 607
- 5765-392 Parallel OSL 605, 617
- 5765-543 Parallel Environment V2.4 605, 612, 618
- 5765-646 recoverable VSD 607
- 5765-B95 GPFS 609
- 5765-C42 Parallel ESSL 605, 614
- 5765-D51 PSSP 205
- 5765-D51 PSSP, Parallel software 605
- 5765-D61 LoadLeveler V2.1.0 605, 611
- 6006 2-meter Drawer-to-Drawer Control Cable 184
- 6007 15-meter Rack-to-Rack Control Cable 184
- 6008 6-meter Rack-to-Rack Control Cable 184
- 6173 Side-Mounted, 3-Phase 195
- 6174 Side-Mounted, 3-Phase, Swiss 195
- 6206 PCI Single-Ended Ultra SCSI Adapter 353,

633, 634
 6206 PCI single-ended Ultra SCSI Adapter 634
 6206 PCI Single-Ended Ultra-SCSI Adapter 633
 6207 PCI Differential Ultra SCSI Adapter 354
 6208 SCSI-2 Fast/Wide Adapter 4-A 355
 6209 SCSI-2 Differential Fast/Wide Adapter 4-B 356
 6215 SSA RAID 5 Adapter with 6222 SSA Fast-Write Cache Option 357
 6222 SSA Fast-Write Cache Option 357
 6227 Gigabit Fibre Channel Adapter 363
 6310 ARTIC960RxD Quad Digital Trunk Adapter 337
 64-bit PCI slots 189
 64-bit System 509
 64-bit technology 13
 6503 Remote I/O Hub - Dual Loop 200
 6913 1000 Watt AC Power Supply 177
 6914 Programmable Power Regulator 177
 6915 Processor Power Regulator 177
 7007 Rackless System Order - Plant Install 204
 7009 Rackless System Order - Field Install 204
 7013-S70, see Model S70 153
 7013-S7A, see Model S7A 154, 155
 7014-S00 Rack 195
 7015-R00 Rack 195
 7015-S70, see Model S70 153
 7015-S7A, see Model S7A 154, 155
 7017-S70, see Model S70 153
 7017-S7A, see Model S7A 154, 155
 7025-F30 73
 7025-F40 73
 7025-F50 73
 7025-F80 85
 7026-H10 73
 7026-H50 73
 7026-H70 73
 7026-H80 99
 7026-M80 136
 7043 Models 25
 7043-140 28
 7043-260 47
 7043-270 53
 7044-170 40
 7044-270 53
 7133-D40 Serial Disk System, Advanced 430
 7133-T40 Serial Disk System, Advanced 430
 7203-001 portable disk drive 449
 7204-118 external SCSI disk drive 448
 7206-005 external 4 mm DAT tape drive 453
 7206-110 external 4 mm DAT tape drive 453
 7207-012 1/4" cartridge tape drive 455
 7207-122 1/4" cartridge tape drive 455
 7207-315 MLR1 tape drive 455
 7210 CD-ROM drive 480
 7331-305 8 mm tape library 477
 7332-005 4 mm DDS tape autoloader 475
 7337 digital linear tape libraries 479
 79 inch tall frame 206
 8 port SP switch 206
 8133 RJ-45 to DB-25 converter cable 171, 203, 210
 8136 128-port Asynchronous cable 171, 203, 210
 8136 rack-mountable remote asynchronous node 171, 203, 210
 8396 SP System Attachment Adapter 171, 205
 8-port Asynchronous Adapter (# 2944) 333
 9077-04S 4-slot SP Switch Router 298
 9077-16S 16-slot SP Switch Router 298
 9171 Side-Mounted, 1-Phase 194
 9173 Side-Mounted, 3-Phase 194
 9174 Side-Mounted, 3-Phase, Swiss 194
 9930 SP SMP node enclosure 270, 281
 9931 SP SMP node enclosure cover plate 281

A

active backplane 162, 182, 184
 adapter capacities 196
 adapter placement guide 661
 adapter placement guidelines
 Model 140 664
 Model 150 667
 Model 170 671
 Model 260 674
 Model 270 677
 Model F50 682
 Model F80 685
 Model H50 687
 Model H70 690
 Model H80 693
 Model M80 696
 Models S70 and S70 Advanced 700
 Adapters
 1000BaseLX media card 300
 1000BaseSX media card 300
 Gigabit Ethernet Media Card 299
 hot-plug adapters servers 504

- hot-plug adapters-WebSM 504
- Hot-Plug in SP Expansion I/O 263
- Hot-plug support 504
- SP High Performance Switch LC-8 Adapter (# 4007) 292
- SP Switch MX Adapter (# 4022) 292, 313
- SP Switch MX2 Adapter (# 4023) 292
- SP Switch-8 Adapter (# 4008) 292
- adapters, communications 329
 - 128-Port Asynchronous Adapter EIA-232 333
 - 8-port Asynchronous Adapter 332
 - ARTIC
 - 2947 ARTIC960Hx 4-Port Selectable Adapter 335
 - 2948 ARTIC960Hx 4-Port T1/E1 Adapter 336
 - 6310 ARTIC960RxD Quad Digital Trunk Adapter 337
- asynchronous
 - 2943 8-Port Asynchronous Adapter 332
 - 2944 128-Port Asynchronous Adapter 333
- ATM
 - 2963 TURBOWAYS 155 PCI UTP ATM Adapter 340
 - 2988 TURBOWAYS 155 PCI MMF ATM Adapter 341
 - 2998 TURBOWAYS 25 ATM Adapter 342
- ESCON
 - 2751 S/390 ESCON Channel Adapter 326
- Ethernet
 - 2968 10/100 Ethernet TX PCI Adapter 342
 - 2969 Gigabit Ethernet-SX PCI Adapter 344
 - 2987 Ethernet AUI/RJ-45 Adapter 346
 - 4951 4-Port 10/100 Base-TX Ethernet Adapter 347
- FDDI
 - 2741 FDDI SK-NET LP SAS Adapter 329
 - 2742 FDDI SK-NET LP DAS Adapter 330
 - 2743 FDDI SK-NET UP DAS Adapter 331
- feature code summary table (PCI) 323
- HIPPI 325
- ISDN
 - 2708 Eicon ISDN DIVA PRO 2.0 S/T Adapter 325
- SP System Attachment Adapter 328
- SysKonnnect SK-NET FDDI-UP SAS 331
- token ring
 - 4959 4/16 Mbps token-ring adapter 348
- adapters, storage 349
 - feature code summary 323
 - Fibre Channel
 - 6227 Gigabit Fibre Channel Adapter 363
 - SCSI
 - 2493 SCSI-2 Fast/Wide RAID Adapter 349
 - 2494 PCI 3-Channel Ultra SCSI RAID Adapter 350
 - 6206 PCI Single-Ended Ultra SCSI Adapter 353
 - 6207 PCI Differential Ultra SCSI Adapter 354
 - 6208 SCSI-2 Fast/Wide Adapter 4-A 355
 - 6209 SCSI-2 Differential Fast/Wide Adapter 4-B 356
 - SSA
 - 6215 SSA RAID 5 Adapter with 6222 SSA Fast-Write Cache Option 357
- additional power components 177
- Address Resolution Protocol 535
- Advanced Interactive Executive 489
- AIX 489, 505, 512
 - 24-Way SMP 496
 - 32-bit to 64-bit 509
 - 4.3.3.0-03 package 490
 - 64-bit
 - execute 64-bit applications 511
 - 64-bit instructions 510
 - 64-bit systems support-servers 509
 - Alternate Disk Installation 500
 - alternative memory allocation (malloc) 496
 - APAR IY09814 490
 - AutoFS 513
 - Bidirectional (BIDI) 522
 - Binary Compatibility 520
 - Binder Library 506
 - C RTE 511
 - C++ V4 RTE 511
 - Code Page 943 522
 - Compatibility Between AIX Versions 3 and 4 556
 - Console Logging 499
 - developing libraries 510
 - device drivers 510
 - directory exploitation 540
 - Documentation 506
 - Base Documentation CD 506
 - Documentation Search 507
 - Extended Documentation CD 506
 - HTML-based 507

- PDF format 507
- Unified Documentation Library Services 508
- Dynamic Processor Deallocation 501
 - Restarting an Aborted Processor Deallocation 503
- Error Messaging 514
- EuroReady 521
- file size maximum 497
- file system size maximum 497
- hardware exceptions 489
- Hardware Requirements 489
- Hardware Support 489
- Hardware support 489
- header files 510
- Hot-plug Adapters 504
- Improved Scheduler Responsiveness 506
- Improved Serviceability Features 499
- International language support 522
- inter-process communication resources 510
- ISO Standard ISO8859-15 520
- JFS
 - Online JFS Backup 513
- kernel 497
 - AutoFS 513
 - dynamically extendable kernel 498
 - I/O 497
 - kernel data addressability 498
 - kernel debugger 498
 - kernel trace buffers 498
 - Process Management 497
- kernel extensions 510
- Korean Input Method 523
- levels supported on SP nodes 309
- Libraries
 - Non-thread-safe libraries 510
 - thread-safe libraries 510
- licenses, unlimited 558
- Logical Volume Manager 511
- logical volumes
 - Concurrent mode 512
 - Online Mirror Backup 512
 - Raw LV online mirror backups 512
 - Super Strict 512
- LP64 511
- maintenance 489, 490
- MAN command 507
- map large files 509
- maximum file system size 497
- Memory Management 496
 - memory/disk requirements 492
- Message Queues 505
- messaging facility 514
- mirroring 511
- mksysb on cd 505
- Multi-Byte Character Set (MBCS) support 523
- multiple threads 510
- Network Installation Manager (NIM) 505
- new hardware
 - booting 490
 - new hardware cd level 489
 - new hardware support 489
 - Model F80, H80.M80.S80,B50 489
- New Locales 524
- NIM
 - replication of resources 505
 - number of disks in a volume group 497
 - number of physical partitions per disk 497
- ODM 503, 505
- open files - limits 497
- physical partitions 497, 512
- print job numbers 556
- Processor State Considerations 503
- Publications 559
- RAID10 511
- Releases and Support Dates 494
- Replaceable Malloc 496
- replacepv 512
- Security
 - SecureWay Directory 540
- Security Certifications 521
- semaphores 505
- Single-Byte Character Set (SBCS) 522
- Standards 519
- Support Mirroring for System Dump 513
- System Management 494
- Taiwanese Input Method 523
- threads per process 497
- Tivoli-Ready 520
- Unicode 522
- Uni-processor 496
- unlimited license 558
- UTF-8 523
- Version 3 and Version 4 Client And Server Compatibility 557
- Volume group 497
- VPN Certification 521
- Web-based System Manager 500

- functions 501
- hardware problem determination 500
- Web-based System Manger
 - HTML-based applications 500
- WebSM 500
 - current resource utilization 501
 - functions 501
 - hardware problem determination 500
 - HTML-based applications 500
- Welcome Center 559
- Workload Manager
 - workload manager 498
 - Workload Manger 496
- XCOFF commands 509
- AIX Bonus Pack 562
 - Database Software 563
 - Development Tools 564
 - e-business & Commerce 568
 - How to order? 570
 - Interoperability Support 567
 - Multimedia 567
 - System Management and Network Management Utilities 565
- AIX Fast Connect 548, 591, 631
 - DCE authentication 548
 - functions 549
 - JFS,CDFS,NFS,CIFS 548
 - logon server 548
 - memory/disk requirements 492
 - Windows 95/98/2000/NT 548
- AIX/ESA 245
- AIX-unlimited user license 558
- ANSI standard
 - X3.131-1986 391
 - X3T9.2/375R Revision 10K 392
- Apache 493
- APAR IY09814 490
- arbiter chip 161
- architectures
 - NUMA-Q 315
 - SCSI-III 392
 - Seascape Storage Architecture 421
 - Serial Storage Architecture (SSA) 423
- Argonne National Laboratories 245
- ARTIC adapters
 - 2947 ARTIC960Hx 4-Port Selectable Adapter 335
 - 2948 ARTIC960Hx 4-port T1/E1 Adapter 336
 - 6310 ARTIC960RxD Quad Digital Trunk Adapter 337
- ASCI 245
- ASCI terminal 203
- ASCII console 751
- asynchronous adapters
 - 2943 8-Port Asynchronous Adapter 332
 - 2944 128-Port Asynchronous Adapter 333
 - 8-port Asynchronous Adapter 332, 333
- asynchronous communication, SCSI 398
- ATM adapters
 - 2963 TURBOWAYS 155 PCI UTP ATM Adapter 340
 - 2988 TURBOWAYS 155 PCI MMF ATM Adapter 341
 - 2998 TURBOWAYS 25 ATM Adapter 342
- ATM LAN emulation 530
- Automount 513
- available graphics accelerators 368
- AWadv-03, benchmark 368

B

- backplane 162, 182, 184
- base power regulators for SP frames 251
- Basic Network Utilities
 - uucp 530
 - uupick 530
 - uusend 530
 - uuto 530
- Basic Network Utilities (BNU) 530
- benchmark data 643
- Binder Library 506
- block diagram explanation 2
- BNC Ethernet 207
- BNU
 - data transfer utility 530
- boot device 185
- boot disk bay 96, 132
- branch buffer 12, 176
- branch unit 12, 176
- bus arbiter chip 161, 180
- bus frequency 182, 197
- bus transfer speed 182

C

C

- 32-bit and 64-bit 574
- 5765-C64 compiler 571
- 64-bit C compiler 510

- advantages 574
 - graphical debugger tool 574
 - LoadLeveler 612
 - SMP 574
 - versions 577
 - withdrawn versions 577
 - C for AIX V4.4 574
 - C Version 5 571
 - C++, LoadLeveler 612
 - cables, SCSI (various types) 634
 - cabling the PCI differential Ultra SCSI adapter 637
 - CacheFS in ONC+ 530
 - Catia 621
 - CATweb 623
 - Operating Environment 622
 - Includes Programs 623
 - CDE 515
 - CDLI 530
 - CDMF 539
 - CD-ROM 418
 - 32X 418
 - CEC cabling 182
 - CIFS Services 539
 - circuit breaker for SP computer room 251
 - cldare 630
 - clinfo 626
 - CLIO/S 327
 - cllockd 626
 - clocking frequency 182
 - clsmuxpd 625, 626
 - clstrmgr 626
 - Cluster Technology, PSSP 608
 - clverify 630
 - coexistence between PSSP levels on the SP 309
 - common command set, SCSI 392
 - Common Desktop Environment 515
 - comparison between SSA and SCSI architecture 404
 - Compilers
 - C
 - Compatibility 573
 - Hardware Requirements 572
 - Redistribution Information 573
 - Software Requirements 572
 - Usage Restriction 573
 - VisualAge C++V4 576
 - VisualAge C++V5 574
 - C for AIX V5 571
 - C Version 4.4 574
 - Fortran
 - XL Fortran 6.1 580
 - XL Fortran 7.1 578
 - XL Fortran RTE 580
 - compilers 571
 - complex fixed point unit 12, 176
 - computational fluid dynamics 246
 - computer system block diagram explanation 2
 - computer-aided design, manufacturing, and engineering software 621
 - Concurrent LVM 626
 - conductor size for SP computer room 251
 - configuration notes
 - HA-H70 121
 - Model 140 34
 - Model 150 39
 - Model 170 45
 - Model 260 52
 - Model 270 59
 - Model F50 84
 - Model F80 98
 - Model H50 109
 - Model H80 133
 - Model M80 146
 - Solution HA-F80 99
 - Solution HA-H80 135
 - Solution HA-M80 150
 - configurator sessions for ordering attached S70/S7A 305
 - console 203, 751
 - consolidated accounting, PSSP 606
 - control workstation 206
 - Control Workstation disk and memory requirements 297
 - Control Workstation for the SP 293
 - copper chip 11, 175
 - copper technology 189
 - Cornell University 245
 - cross-port traffic 182
 - C-SPOC 589, 628, 630
 - customer-installable features and models 753
 - CWS 206
- D**
- DARE 628
 - DAT (Digital Audio Tape)
 - 7332-005 4mm DDS tape autoloader 475
 - DAT tape drives 453

- data buffers 180
- data traffic redundancy 190
- Dataserver 7051
 - AIX support 489
- DB2 540
- DDR 11, 175
- DES 539
- design overview, RS/6000 2
- Deskside Server 69
 - description 69
 - F50 69
 - F80 85
- device-to-device cables 638
- DHCP 538
- differential, SCSI 398
- disk drives, external
 - 7203 portable disk drive 449
 - 7204 external disk drives
- disk systems, external 424
 - 2102-F10 Fibre Channel RAID Storage Server 427
 - 7133-D40 Serial Disk System, Advanced 430
 - 7133-T40 Serial Disk System, Advanced 430
 - Expandable Storage Plus 439
 - product positioning 424
 - summary table 425
- dispatch buffer 12, 176
- display panel 196
- Displays
 - 6091 491
 - G52 491
 - P200 491
 - P201 491
 - P50 491
 - P70 491
- DLPI 530
- DLT tape drive
 - 7205-311 digital linear tape drive 451
- Double Data Rate 11, 175
- DRV-06, benchmark 368
- DX-05, benchmark 368
- Dynamic CPU Deallocation 501
- Dynamic Load Name Resolver APIs 532
- Dynamic Processor Deallocation 501

E

- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter 325
- electrical current for SP power supplies 251

- electrical power requirements for SP frames 251
- en0 206
- engineering analysis 246
- engineering design 246
- Enterprise Servers
 - Model M80 136
 - Model S70 153
 - Model S7A 154
- Entry-Level Enterprise Server 69
- EPOW
 - Model F80 EPOW 89
- ESCON 244
- ESCON adapters
 - 2751 S/390 ESCON Channel Adapter 326
- ESCON Multiple Image Facility 327
- ESCRM 590, 631
- EtherChannel 537
- Ethernet adapters 206
 - 2968 10/100 Ethernet TX PCI Adapter 342
 - 2969 Gigabit Ethernet-SX PCI Adapter 344
 - 2987 Ethernet AUI/RJ-45 Adapter 346
 - 4951 4-Port 10/100 Base-TX Ethernet 347
- event emulation 630
- execution rate 12, 176
- external device support table, all RS/6000's 747
- external disk drives
 - see disk drives, external
- external disk storage 424
 - see also storage, external
- external node 207
- external storage 421
 - see also storage, external
- extranet 246

F

- F50, upgrade paths F80 98
- Facts and Features 17
- fans 196
- Fast Connect/memory/disk requirements 492
- fast SCSI 397
- Fast Trap signal 514
- fast/wide, regarding SCSI-2 397
- fast-20 wide
 - SCSI 397
- FDDI 331
- FDDI adapters
 - 2741 FDDI SK-NET LP SAS Adapter 329
 - 2742 FDDI SK-NET LP DAS Adapter 330

2743 FDDI SK-NET UP DAS Adapter 331
 SysKonnnect SK-NET FDDI-LP SAS PCI 329
 SysKonnnect SK-NET FDDI-UP SAS adapter 331
 feature codes
 2050 SP 332 MHz SMP Thin node 281
 2051 SP 332 MHz SMP Wide node 270
 2104 LVD expandable storage plus 439
 2493 SCSI-2 Fast/Wide RAID Adapter 349
 2494 PCI 3-Channel Ultra SCSI RAID Adapter 350
 2708 Eicon ISDN DIVA PRO 2.0 S/T Adapter 325
 2741 FDDI SK-NET LP SAS Adapter 329
 2742 FDDI SK-NET LP DAS Adapter 330
 2743 FDDI SK-NET UP DAS Adapter 331
 2751 326
 2943 8-Port Asynchronous Adapter 332
 2944 128-Port Asynchronous Adapter 333
 2947 ARTIC960Hx 4-port Selectable Adapter 335
 2948 ARTIC960Hx 4-Port T1/E1 Adapter 336
 2962 2-port multiprotocol X.25 Adapter 338
 2963 ATM 155 Turboways UTP Adapter 340
 2968 10/100 Ethernet TX PCI Adapter 342
 2969 Gigabit Ethernet-SX PCI Adapter 344
 2987 Ethernet AUI/RJ-45 Adapter 346
 2988 TURBOWAYS 155 PCI MMF ATM Adapter 340, 341
 2998 Turboways 25 ATM Adapter 342
 4007 SP High Performance Switch LC-8 Adapter 292
 4008 SP Switch-8 Adapter 292
 4010 High Performance Switch 292
 4011 SP Switch 292
 4022 SP Switch MX Adapter 292, 313
 4023 SP Switch MX2 Adapter 292
 4951 4-Port 10/100 Base-TX Ethernet PCI Adapter 347
 4959 4/16 Mbps token-ring adapter 348
 6206 PCI Single-Ended Ultra SCSI Adapter 353
 6206 PCI single-ended Ultra SCSI Adapter 633, 634
 6206 PCI single-ended Ultra SCSI adapter 634
 6207 PCI Differential Ultra SCSI Adapter 354
 6208 SCSI-2 Fast/Wide Adapter 4-A 355
 6209 SCSI-2 Differential Fast/Wide Adapter 4-B 356
 6215 SSA RAID 5 Adapter 357
 6222 SSA Fast-Write Cache Option 357
 6227 Gigabit Fibre Channel Adapter 363
 6310 ARTIC960RxD Quad Digital Trunk Adapter 337
 8396 SP System Attachment Adapter 171, 205
 9930 SP SMP node enclosure 270, 281
 9931 SP SMP node enclosure cover plate 281
 features
 optional
 supported on Model 140 30
 supported on Model 150 36
 supported on Model 170 42
 supported on Model 260 48
 supported on Model 270 55
 supported on Model F50 78
 supported on Model F80 90
 supported on Model H50 105
 supported on Model H70 113
 supported on Model H80 127
 supported on Model M80 142
 supported on Model S7A 167
 standard
 on Model 140 28
 on Model 150 35, 61
 on Model 170 41
 on Model 260 47
 on Model 270 53
 on Model F50 74
 on Model F80 85
 on Model H50 101
 on Model H70 111
 on Model H80 123
 on Model M80 141
 on NUMA-Q 2000 Model 410 318
 Fibre Channel adapters
 6227 Gigabit Fibre Channel Adapter 363
 file size maximum 497
 file system size maximum 497
 firmware 185
 fixed point unit 12, 176
 floating point unit 12, 176
 fork interface 540
 Fortran 578
 HPF 580
 LoadLeveler 612
 LPEX 581
 MPI 580
 versions 581

- xldb 581
- xxlf 581
- Frame Contents 247
- frame descriptions for the SP 247
- frame dimensions for the SP 250
- frame limits 207
- fsck 630

G

- General Parallel File System, see GPFS 605, 609
- general SCSI consideration 400
- Gigabit Ethernet adapter 344
- GL 3.2
 - supported hardware 490
- GPC/OPC, benchmark 368
- GPC/PLB, benchmark 368
- GPC/XPC, benchmark 368
- GPFS 605, 609
 - administration 610
 - standards compliance 610
- Graphics 514
- graphics accelerators
 - available graphics accelerators 368
 - hardware classes 367
 - withdrawn graphics accelerators 387
- Graphics adapters 367
 - POWER Gt3 491
 - POWER Gt3i 491
 - POWER Gt4 491
 - POWER Gt4e 491
 - POWER Gt4i 491
 - POWER Gt4x 491
 - POWER Gt4xi 491
 - POWER GXT1000 491
 - POWER GXT150 491
 - POWER GXT150L 491
 - POWER GXT150M 491
 - POWER GXT150P 491
 - POWER GXT155L 491
 - POWER GXT2000P 491
 - POWER GXT250P 491
 - POWER GXT255P 491
 - POWER GXT3000P 491
 - POWER GXT500 491
 - POWER GXT500D 491
 - POWER GXT500P 491
 - POWER GXT550 491
 - POWER GXT800M 491

- POWER GXT800P 491
- POWER GXT800PT 491
- graphics unsupported on SP 314
- graPHIGS 519
 - default swapping 519
 - Improved clipping performance 519
 - maximum data addressability 519
 - memory/disk requirements 492
 - performance enhancements 519
 - supported hardware 490
- ground cable 207
- groupware application 246
- GUIs 514
- GXT120P, graphics accelerator 384
 - display and cable 385
 - supported machines 384
- GXT130P, graphics accelerator 381
 - display and cable 382
 - performance table 383
 - supported machines 381
- GXT2000 517
- GXT2000P, graphics accelerator 373
 - display and cable 374
 - performance table 375
 - supported machines 373
- GXT250P, graphics accelerator 379
 - display and cable 380
 - supported machines 379
- GXT255P, graphics accelerator 379
 - display and cable 380
 - supported machines 379
- GXT3000 517
- GXT3000P, graphics accelerator 369
 - display and cable 371
 - performance table 372
 - supported machines 370
- GXT300P, graphics accelerator 376
 - display and cable 377
 - performance table 378
 - supported machines 377

H

- H70, upgrade paths H80 134
- HACMP 4.4 592
- HACMP and HACMP/ES Version 4.2.2 630
- HACMP and HACMP/ES Version 4.3 630
- HACMP and HACMP/ES Version 4.3.1 631
- HACMP and HACMP/ES Version 4.4 632

- HACMP Version 1.1 625
- HACMP Version 2.1 626
- HACMP Version 3.1 626
- HACMP version 3.1.1 627
- HACMP version 4.1 627
- HACMP version 4.1.1 627
- HACMP version 4.2.0 628
- HACMP Version 4.2.1 628
- HACMP/ES Version 4.2.1 629
- HACWS, PSSP 607
- HAGEO 628
- hardware components for the SP 246
- hardware control and supervision for the SP 249
- HA-S70 Advanced Cluster Server, see solution
- HA-S70 208
- HA-S7A 153
- HA-S80 153
- Hashed Message Authentication Code 534
- High Availability 294, 583
 - Cluster Multi-Processing, HACMP 584
 - Geographic Cluster Product, HAGEO 597
 - key factors 583
 - Web site 601
- high availability Control Workstation for the SP 294
- High Performance Switch (# 4010) 292
- HIPPI adapters 325, 326
- history
 - of RS/6000 product line 1
 - SP 244
- hot-plug PCI slots 88
- hot-pluggable PCI adapters 504
- HPS 627
- HPS switch 206
- HTTP GET Engine 537
- HTTP Server 493

I

- I/O bus
 - SP 332 MHz SMP Thin node 281
 - SP 332 MHz SMP Wide node 271
- IBM 7135-110 626
- IBM 9333 626
- IBM C for AIX V5 571
- IBM Long-wave Serial HIPPI PCI Adapter 326
- InfoWindow 751
- installation by customer (features and models) 753
- instruction buffer 12, 176
- instructions per cycle 12, 176

- Internet 246
- intranet 246
- investment protection 156
- IP Addresses - Max. number per adapter 537
- IP Security on IPv4 and IPv6 533
- IPv6 531
 - Supported Standards 532
- IPX/SPX protocols 530
- Irving Wladawsky-Berger 244
- ISDN adapters 325
 - 2708 Eicon ISDN DIVA PRO 2.0 S/T Adapter 325

J

- J30, upgrade paths S7A 154
- J40, upgrade paths S7A 154
- J50, upgrade paths S7A 154
- Japan Kit Fixed Fonts 524
- Java 550
 - Big Decimal 552
 - Improved Java performance 539
 - Java Development Toolkit (JDK) 550
 - Java Port 550
 - Just-In-Time (JIT) Compiler 550
 - Swing 552
- Java Database Connectivities 551
- Java Remote Method Invocation 551
- Java Security 551
- JFS locking 539

K

- kerberos authentication 628
- Kerberos primary server 293
- Kerberos secondary server 293
- kernel
 - kernel extension
 - HTTP Get Engine 537

L

- L1 cache latency 11, 175
- L2 cache 11, 175
- L2 cache bandwidth 11, 175
- LAN adapters
 - 2741 FDDI SK-NET LP SAS Adapter 329
 - 2742 FDDI SK-NET LP DAS Adapter 330
 - 2743 FDDI SK-NET UP DAS Adapter 331
 - 2963 TURBOWAYS 155 PCI UTP ATM Adapter

- 340
- 2968 10/100 Ethernet TX PCI Adapter 342
- 2969 Gigabit Ethernet-SX Adapter 344
- 2987 Ethernet AUJ/RJ-45 Adapter 346
- 2988 TURBOWAYS 155 PCI MMF ATM Adapter 341
- 2998 TURBOWAYS 25 ATM Adapter 342
- 4951 4-Port 10/100 Base-TX Ethernet Adapter 347
- 4959 4/16 Mbps token-ring adapter 348
- 6227 Gigabit Fibre Channel Adapter 363
- LAPI, Parallel software 608
- Large Scale Servers 243
- latency 10, 175, 181
- LDAP 527, 540
 - DSSeries LDAP Directory 527
 - eNetwork LDAP Directory 528
 - LDAP Software Requirements 529
 - LDAP Standards 529
 - libraries 544
 - Secure Sockets Layer (SSL) 529
- legacy models for the SP 248
- libbsd.a 510
- libc.a 510
- libdes.a 510
- libm.a 510
- libmsaa.a 510
- libs.a 510
- License Use Manager (LUM) 506
- Light-03, benchmark 368
- LINPACK
 - benchmark data 652
 - benchmark definition 659
- load/store unit 12, 176
- LoadLeveler 605, 611
 - central control 612
 - highlights 611
 - individual control 612
 - interfaces 611
 - scalability 612
 - Web site 612
 - what to expect 612
- logical frame 207
- logical nodes 207
- Logical Volumes 508
- Lotus Domino Go Server 493
- Lotus Domino Server 493
- LowBoy frame 206
- LP, Linear Programming 618

M

- MA-M80 150
- man commands 507
- maximum file size 510
- maximum file system size 510
- maximum RIO loop 188
- maximum SPCN loop 183
- medical simulation 246
- memory bandwidth 180
- memory books 182
- memory connectors 177
- memory conversion 201
- memory quad physical layout 179
- memory redundancy 177
- memory scrubbing 177, 189
- Message Queues 505
- MIB 625
- microprocessors
 - architectures 7
 - see also processors
- Microsoft Internet Information Server 493
- Midrange Enterprise Server
 - definition 69
- Mid-range Servers
 - Model H80 122
- migratepv 512
- migration, PSSP 607
- minimum RIO loop 185
- minimum SPCN loop 183
- MIP, Mixed-Integer Programming 618
- Mirror Resynchronization 512
- mirroring 512, 513
 - Support Mirroring for System Dump 513
- mission critical application 625
- mixing processor types 174
- mksysb on cd 505
- Mode 1 625
- Mode 2 625
- Model
 - F50 69
 - H50 70
 - H70 70
 - HA-H70 71
 - M80 136
- Model 140
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 664
 - configuration notes 34

- customer-installable features 753
- external device support table 748, 749
- facts and features summary 18
- history and use 26
- optional features 30
- performance data 643
- picture 28
- product description 28
- publications 34
- site and hardware planning information 711
- standard features 28
- upgrade paths 27
- Model 150
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 667
 - configuration notes 39
 - customer-installable features 753
 - external device support table 748, 749
 - facts and features summary 18
 - history and use 26
 - optional features 36
 - performance data 643
 - picture 35
 - publications 40, 66
 - site and hardware planning information 713
 - standard features 35, 61
 - upgrade paths 27
- Model 170
 - adapter placement guidelines 671
 - configuration notes 45
 - history and use 26
 - optional features 42
 - product description 40
 - publications 46
 - site and hardware planning information 714
 - standard features 41
- Model 260
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 674
 - AIX support 490
 - configuration notes 52
 - customer-installable features 753
 - external device support table 748, 749
 - facts and features summary 18
 - history and use 26
 - optional features 48
 - performance data 643
- picture 47
- product description 47
- publications 52
- site and hardware planning information 716
- standard features 47
- upgrade paths 27
- Model 270
 - adapter placement guidelines 677
 - configuration notes 59
 - history and use 27
 - optional features 55
 - product description 53
 - publications 60
 - site and hardware planning information 718
 - standard features 53
- Model 7014-S00 119
- Model 730
 - AIX Support 489
- Model 822
 - AIX support 489
- Model 823
 - AIX support 489
- Model 851
 - AIX support 489
- Model 860
 - AIX support 489
- Model B50
 - AIX support 489
 - site and hardware planning information 719
- model conversion 199
- Model F30 73
- Model F40 73
 - adapter descriptions, communications and storage 323
 - customer-installable features 753
 - external device support table 748, 749
 - facts and features summary 19, 20
 - performance data 643
- Model F50
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 682
 - configuration notes 84
 - customer-installable features 753
 - disk bays optional 83
 - external device support table 748, 749
 - F50 69
 - facts and features summary 19, 20
 - features 74

- memory 83
- memory additional 82
- optional features 77
- performance data 643
- processor optional 82
- processors 83
- publications 84
- service director 77
- Service Processor 77
- site and hardware planning information 720
- standard features 74
- upgrade paths 73
- Model F80 85
 - 6-pack 96
 - adapter placement guidelines 685
 - audio features 97
 - Auto-dial-out 89
 - boot disk bay 96
 - configuration notes 98
 - disk bays 96
 - disk drives internal 96
 - features 90
 - front view 96
 - graphics accelerators 97
 - hot-plug PCI slots 88
 - hot-swappable disks 88, 96
 - I/O drawer 124
 - media bay 97
 - media bays 96
 - memory 87
 - memory additional 95
 - memory board 87, 95
 - multimedia features 97
 - optional features 89
 - power supplies redundant 98
 - processor optional 95
 - processor subsystem 87
 - publications 99
 - quads 87, 95
 - RAID 88, 96
 - rear view 87
 - redundant power and cooling 88
 - reliability, availability, serviceability features 88, 89
 - required AIX 489
 - Service Director 88
 - service director 88
 - Service Processor 89
 - SSA and SSA RAID 97
 - SSA disks 96
 - standard features 85
 - tape drives internal 97
 - upgrading from Model F50 95, 98
- Model H10 73
- Model H50 73
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 687
 - bays 104, 108
 - configuration notes 109
 - customer-installable features 753
 - disk internal 108
 - facts and features summary 21
 - features 101
 - graphics accelerator 109
 - HIPPI adapter features 106
 - memory additional 103
 - optional features 104
 - performance data 643
 - power supply AC redundant 109
 - publications 109
 - SCSI adapters 108
 - SCSI RAID 108
 - site and hardware planning information 721
 - standard features 101
- Model H70 119
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 690
 - bays 117
 - customer-installable features 753
 - facts and features summary 21
 - features 110
 - fibre channel adapter features 115
 - HIPPI adapter features 115
 - memory 117
 - optional features 112
 - performance data 643
 - power supply redundant 119
 - processors 112, 117
 - RAS 119
 - site and hardware planning information 723, 724, 727
 - slot assignments 119
 - SSA 118
 - standard features 111
 - upgrade paths 73
- Model H80

- adapter placement guidelines 693
- AIX support 489
- ARTIC adapter features 129
- async adapter features 130
- ATM adapter features 130
- boot disk bay 132
- Central Electronics Complex, Model H80
 - CEC 122
- configuration notes 133
- conversion to Model M80 139
- DC power supplies 133
- debug port 126
- diagnostics port 126
- digital trunk adapter features 130
- disk internal 132
- diskette drive 122
- Ethernet adapter features 130
- FDDI adapter features 131
- features 122
- fibre channel adapter features 129
- graphics Accelerator features 127
- graphics accelerators 133
- HIPPI adapter features 130
- hot-plug PCI slots 125
- I/O drawer 122
- I/O drawer inner view 125
- I/O drawer rear view 126
- internal CD-ROM features 129
- internal disk bay features 128
- internal disk features 128
- internal tape drive features 128
- IPL disks 132
- ISDN adapter features 131
- JTAG cables 125
- media bay 125
- media bays 132
- memory 124
- memory additional 131
- memory board 124
- memory features 127
- miscellaneous features 131
- optional features 126
- power supplies redundant 133
- power supply features 127
- processor features 127
- processor optional 131
- processor subsystem 124
- publications 134
- quads 124, 131

- racks 126
- reliability, availability, serviceability features 124
- remote I/O cables 124
- required AIX 489
- RIO cables 124
- SCSI adapter features 129
- secondary I/O drawer 125
- Service Processor 124
- SPCN cables 125
- SSA adapter features 129
- system adapter features 130
- token ring adapter features 130
- Upgrading from Model H70 134
- UPS 133
- V/S COMM cable 125
- WAN adapter features 130
- Model HA-F80
 - Solution HA-F80
 - HA-F80 99
- Model HA-H70 120
- Model HA-H80
 - Solution HA-H80 135
- Model HA-M80 150
- Model M80 136
 - adapter placement guidelines 696
 - AIX support 489
 - configuration notes 146
 - HIPPI adapter features 145
 - history and use 136
 - memory 88
 - optional features 142
 - picture 137
 - product description 137
 - publications 150
 - required AIX 489
 - standard features 141
- Model N40
 - AIX support 489
- Model S70 153
 - adapter descriptions, communications and storage 323
 - adapter placement guidelines 700
 - configurator sessions for attaching to SP 305
 - customer-installable features 753
 - performance data 643
 - site and hardware planning information 729, 734
- Model S70 Advanced

- adapter descriptions, communications and storage 323
- adapter placement guidelines 700
- customer-installable features 753
- L1 cache 159
- L2 cache bandwidth 159
- memory controller complex 160
- memory ports 160
- performance data 643
- rackless order 203
- RIO cables 184
- serial ports 203
- site and hardware planning information 729, 734
- system bandwidth 162
- system bus 160
- system bus speed 162
- system switch complex 161
- upgrading 199
- Model S70 Advanced Configuration Notes 171
- Model S70 Advanced Overview 157
- Model S70 Advanced Product Description 157
- Model S70 Advanced Publications 172
- Model S70 Advanced Standard Configuration 166
- Model S7A 154, 155
 - configurator sessions for attaching to SP 305
 - Dynamic Processor Deallocation 158
 - I/O drawers 162
 - I/O rack 162
 - memory 160
 - optional features 142, 167
 - pre-install meeting for SP attachment 304
 - Processor Card 160
 - processors 158
 - RIO cables 165
 - Service Director 158
 - Service Processor 158
 - SP attachment 171
 - SPCN 165
 - upgrade paths 154
- Model S80 155, 178
 - AIX support 489
 - announce date 155
 - base configuration 190
 - Cabling 182
 - CD-ROM 156
 - CEC layout 178
 - data flow switches 179
 - Double Data Rate 11, 175
 - Dynamic Processor Deallocation 199
 - maximizing memory performance 178
 - maximum configuration 191
 - Memory 177
 - memory bandwidth 180
 - Memory Controller Complex 179
 - memory controller complex 179
 - memory feature sizes 177
 - memory quad physical layout 179
 - memory quads 180
 - memory redundancy 177
 - memory scrubbing 177
 - optional features 191
 - power supplies 177
 - Power Supply 177
 - pre-install meeting for SP attachment 304
 - Processor 174
 - processor regulator 177
 - publications 199, 205
 - rackless order 203
 - Remote I/O Cables 184
 - RIO cables 184, 188
 - RIO hub 179
 - RS64 III processor card 176
 - S7A to S80 Processor Conversion 200
 - serial ports 203
 - Service Processor 198
 - Service Tools 198
 - SPCN cables 183
 - system bandwidth 180
 - system buses 11, 174, 175
 - System combinations S80/H70 217
 - system switch complex 181
 - upgrade paths 157
- Model S80 Product Description 174
- Models 821
 - AIX support 489
- Motif 2.1 516
 - 32-bit and 64-bit libraries 516
 - Thread-Safe 516
- MPI, Parallel software 608
- MPOA 539
- multiple PDUs 197

N

- N+1 power supplies 250
- native serial ports 203
- NCS 530

- Net View 628
- NetBIOS-over-IP 548
- Netscape Enterprise Server 493
- Netscape FastTrack Server 493
- Network File Cache 536
- Network Installation Manager 505
- Network Storage Manager 486
- NFS 530
- NFS Version 3 630
- NFS3 538
- NIM 505
- NIM master process handling 505
- NIS 526
 - Differences between NIS & NIS+ 526
- NIS+ 525
- node selection for the SP 254
- noise emission notes 742
- NUMA-Q 315
 - architecture 315
 - NUMA-Q 2000 server 318
 - products and solutions 317
- NUMA-Q architecture 315
- NUMA-Q nodes
 - NUMA-Q 2000 Model 410 318
- NUMA-Q Technology 315

O

- ODM 505
- ONC+ 513
- on-chip data buses 13, 177
- On-Demand tunneling 545
- OpenGL
 - 64-bit indirect rendering 517
 - Clip Volume Hint 517
 - Color Blend Extension 518
 - GLX support 517
 - memory/disks requirements 492
 - MultiDrawArray Extension 518
 - Multi-Mode Vertex Arrays 517
 - OpenGL 1.1 and 1.2 517
 - OpenGL Applications Thread Enablement 517
 - performance enhancements 517
 - POWER GXT3000P 518
 - supported hardware 490
 - Texture Mirrored Repeat Extension 518
 - Three new extensions 518
 - Vertex Array Draw Element Range function 518
 - Vertex Array List 517

- ZAPdb Interactive OpenGL Debugger 518
- optical storage devices 480
 - 3995 optical library 482
 - 7210 optical CD-ROM drive 480
- optional features
 - Model 140 30
 - Model 150 36
 - Model 170 42
 - Model 260 48
 - Model 270 55

P

- Paging Space 492
- Parallel Engineering and Scientific Subroutine Library see Parallel ESSL 614
- Parallel Environment 612
 - advantages 613
 - highlights 613
 - key elements 613
 - PE functions 614
 - software requirements 614
- Parallel ESSL 614
 - advantages 615
 - application environments 616
 - compatibility 616
 - highlights 615
 - mathematical functions 615
 - SMP support 616
 - Web sites 617
- Parallel Optimization Subroutine Library, see Parallel OSL 617
- Parallel OSL 617
 - hardware and software prerequisites 619
 - Linear Programming 618
 - Mixed-Integer Programming 618
 - parallel application programming 618
 - performance 619
 - Quadratic Programming 619
 - solution capabilities 619
- parallel processing 245
- Parallel software 605
 - GPFS 605, 609
 - LoadLeveler 611
 - PSSP 605
 - virtual shared disk 608
- Parallel System Support Programs for AIX, see PSSP 605
- Parallel Systems Support Program 155

- Path Maximum Transmission Unit 535
- PC Interoperability 548
- PCI 2.1 189
- PCI 3-Channel Ultra2 SCSI RAID Adapter 323
- PCI adapters
 - see adapters, communications or adapters, storage 323
- PCI bus speed 197
- PCI signaling 189
- PCI slot availability 196
- PCI throughput 189
- PCI-to-PCI bridges 189
- PDU 195, 197
- PDU's
 - SP frames 252
- PE, see Parallel Environment 605, 612
- Performance 199
 - benchmark data for RS/6000 models 643
 - benchmark definitions 653
 - bf 546
 - fdpr 546, 547
 - filemon 546
 - fileplace 546
 - gennames 547
 - ipfilter 547
 - lockstat 546
 - netpmon 546
 - parallel OSL 619
 - pprof 547
 - rmss 546
 - stem 546
 - svmon 547
 - topas 547
 - tprof 547
- Performance Analysis Tools Enhancements 546
- Performance Toolbox Agent (PTX) 547
- peripheral devices for the SP 313
- physical layout 178
- pictures
 - CEC and I/O rack, Model S7A 157
 - I/O drawer, Model S7A 164
- PIOFS, Parallel software 609
- pipeline execution 12, 176
- pipeline mode 189
- PLBsurf93, benchmark 368
- PLBwire93, benchmark 368
- POWER architecture 244
- power distribution unit (PDU) for SP frames 252
- Power Distribution Units
 - Model S7A 164
- power distribution units 195
- power output for SP frames 251
- power supplies
 - Model S80 177
- power supply phases for SP frames 251
- power system upgrade for SP frames 271, 281
- POWER3 Microprocessor 7
- POWER3 SMP High Node 261
- POWER3-II Microprocessor 7
- PPP 530
- Predictive Failure Analysis 189
- pre-install meeting for SP attached servers 304
- primary drawer connection 184
- Printer Support 552
- Pro-CDRS-02, benchmark 368
- processor books 182
- processor regulator 177
- processor upgrade 200
- processors
 - architectures 7
 - key new technologies 14
 - POWER3 7
 - POWER3-II 7
 - RS64/RS64 II 9
- product descriptions
 - Model 140 28
 - Model 170 40
 - Model 260 47
 - Model 270 53
- protocols
 - user space 309
- PSSP 155, 205, 629
 - application programming models 608
 - availability/recoverability 610
 - cluster technology 606, 608
 - collecting performance data 609
 - consolidated accounting 606
 - enabling parallel processing 608
 - features and benefits 606
 - HACWS 607
 - hardware and software prerequisites 607
 - highlights 605
 - levels supported on SP nodes 309
 - licence 207
 - migration 607
 - node isolation 607
 - parallel software 605
 - software packages 609

- subsystem communications support 608
- system monitoring and control 606
- system partitioning 606
- virtual shared disk 607
- PSSP, parallel software 605
- Pthread Debug Library 514
- publications
 - Model 140 34
 - Model 150 40, 66
 - Model 170 46
 - Model 260 52
 - Model 270 60

Q

- quads 87, 124
- Quality of Service 508, 536

R

- R30, upgrade paths S7A 154
- R40, upgrade paths S7A 154
- R50, upgrade paths S70 154
- R50, upgrade paths S7A 154
- rack 119
- rack contents specify code 204
- rack door removal 195
- rack mountable tape units 195
- rack ordering rules 195
- rack power 195
- rackless order 203
- Rack-Mounted Server 70
 - H50 - Entry-Level Enterprise Server 69
 - H50-Entry-Level Enterprise Server 70
 - H70 - Entry-Level Enterprise Server 69, 70
 - HA-H70 71
 - HA-H70 - Entry-Level Enterprise Server Solution 69
- Rack-mounted Server
 - NUMA-Q 2000 Enterprise server 318
- Rack-Mounted-Servers
 - HA-M80 Enterprise Servers 150
 - M80-Enterprise Server 136
- RAID
 - RAID0 511
 - RAID1 511
 - striping 511
- RAN 333
- RAS
 - Model H70 119

- receptacle for SP computer room 251
- reducing footprint 203
- redundant cooling 196
- redundant data paths 184
- redundant power 195, 196
- redundant power supplies for the SP 250
- reliability, availability, serviceability features 76, 102
- Remote Asynchronous Node 333
- Remote I/O cables
 - Model M80 147
 - Model S7A 165
- Remote Reboot 506
- rendering 516
- repeaters, SCSI 399
- Resource groups 627
- RIO 194
- RIO bandwidth 184
- RIO cables
 - Model M80 147
 - Model S7A 165
- RIO cabling 184
- RIO connections 184
- RIO loops 184
- ROLTP
 - benchmark data for SMP systems 647
 - benchmark data for uni-processor systems 643
 - benchmark definition 657
- RS/6000 Model 140 28
- RS/6000 Model 170 40
- RS/6000 Model 260 47
- RS/6000 Model 270 53
- RS/6000 model equivalents for SP nodes 310
- RS/6000 Model F50 74
- RS/6000 Model F80 85
- RS/6000 Model H50 101, 122
- RS/6000 Model H70 110
- RS/6000 Model S70, see Model S70 153
- RS/6000 Model S7A, see Model S7A 154, 155
- RS/6000 SP 155, 243
- RS-232 cables 207
- RS-232 ports 206
- RS64 II processor 159
- RS64 III processor 174
- RS64 III processor card 11, 174, 175
- RS64/RS64 II Microprocessor 9
- RSCT 590, 631
- Russell-Stol 251

S

- S/390 ESCON Channel Adapter 326
- S70, see Model S70 153
- S7A
 - SP attachment 328
- S7A, see Model S7A 154, 155
- S80
 - SP attachment 328
- scalability, LoadLeveler 612
- Scalable Electrical Power Base Unit (SEPBU) for SP frames 252
- scheduler 506
- Schema update operations 541
- SCSI
 - see also Communications and Storage I/O Adapters chapter 398
 - asynchronous 398
 - common command set 392
 - differential 398
 - differential cable lengths 399
 - repeaters 399
 - specifications summary 400
 - summary of specifications 400
 - synchronous 398
 - table of modes/rates 400
 - tagged command queuing 392
 - ultra 397
 - Ultra-2 398
- SCSI adapters
 - 2493 SCSI-2 Fast/Wide RAID Adapter 349
 - 2494 PCI 3-Channel Ultra SCSI RAID Adapter 350
 - 6206 PCI Single-Ended Ultra SCSI Adapter 353
 - 6207 PCI Differential Ultra SCSI Adapter 354
 - 6208 SCSI-2 Fast/Wide Adapter 4-A 355
 - 6209 SCSI-2 Differential Fast/Wide Adapter 4-B 356
- SCSI-2 F/W RAID 323
- SCSI-III 392
- SDRAM modules 177
- Seascope Storage Enterprise Architecture 421
- SecureWay 508
- SecureWay Directory
 - attribute-type definitions 541
 - client access 543
 - migration utility 542
 - Web administration 543
- SecureWay Directory Client SDK 544
- SecureWay Directory Security Enabler 544
- Secureway Directory Server 493
- Security 540
 - Network Security 545
 - Digital Signatures 546
 - Dynamic tunnels 545
 - IKE Tunnels 545
 - Remote TCP/IP Commands 545
 - Tracing IP Security and IKE Messages 545
- seismic processing 246
- Semaphores 505
- Sendmail 526
- SEPBU 252
- serial cables 207
- serial number, Model S7A 154
- serial port connections 206
- serial ports on SP nodes 314
- Service Director 88, 198
- service processor 196, 203
- service processor firmware 206
- Shared Memory Regions 505
- shared-nothing architecture 313
- short frame 248, 250
- Short-wave Serial HIPPI PCI Adapter 325
- single frame system 207
- site and hardware planning guide 711
- SLIP 530
- SMIT 508
- SMP 627
- SNA Alert architecture 514
- SNMP 625
- SNMP Agent 530
- soft memory errors 177
- software requirements, Parallel Environment 614
- Solution HA-H70
 - configuration notes 121
 - features 120
 - standard features 120
- Solution HA-H80
 - standard features 135
- Solution HA-M80
 - configuration 150
 - configuration notes 150
- Solution HA-S70 208
 - configurations 209
 - enhancements 211
 - value proposition 208
- Solution HA-S80 211
 - High-Availability Hardware 213

- High-Availability Software 214
- Maximum Configurations 216
- Solution Highlights 213
- Solutions HA-S80
 - Standard Configuration 212
- SP 243
 - 200 MHz nodes 257
 - 332 MHz SMP Thin node at a glance 258
 - 332 MHz SMP Thin node details 281
 - 332 MHz SMP Thin Node Switch Restrictions 285
 - 332 MHz SMP Wide node at a glance 258
 - 332 MHz SMP Wide node details 270
 - 375 MHz High node 256
 - 375 MHz Nodes 255
 - 375 MHz POWER3-II Wide 265
 - adapter descriptions, communications and storage 323
 - Adapters
 - I/O Adapters Supported on PCI SMP Nodes 307
 - Advantages 300
 - Application Isolation 312
 - coexistence between PSSP levels 309
 - Contents of the Frame 247
 - Control Workstation 293
 - Control Workstation disk and memory requirements 297
 - Disk Pairs - RAID 311
 - documentation 246
 - Dynamic Processor Deallocation 280
 - electrical power requirements for frames 251
 - Expansion I/O Drawer 263
 - Expansion I/O Unit
 - Hot-Plug Adapters 263
 - external node attachment 328
 - F/C
 - 1550 263
 - 2056 275
 - 2057 265
 - 9122 306
 - 9123 306
 - 9222 306
 - 9223 306
 - 9310 306
 - Frame Contents 247
 - frame descriptions 247
 - frame dimensions 250
 - graphics unsupported 314
 - hardware components 246
 - hardware control and supervision 249
 - High Availability 311
 - high availability Control Workstation 294
 - High Performance LC-8 Switch Adapter (# 4007) 292
 - High Performance Supercomputer Systems Development Laboratory 244
 - High Performance Switch (# 4010) 292
 - legacy models 248
 - legacy models by frame and switch types 248
 - legacy models by node types 249
 - Model 500 248
 - Model 550 248
 - Node f/c 2058 259
 - node selection 254
 - nodes 253
 - 375 MHz POWER3-II thin 275
 - Capacity 254
 - nodes at a glance 255
 - origins 244
 - parallel processing 245
 - PdUs for frames 252
 - peripheral devices 313
 - Power Cords 252
 - power supplies 250
 - POWER3 SMP High 375 MHz 256
 - POWER3 SMP High Node
 - floating-point and fixed point execution 261
 - memory bandwidth 261
 - POWER3 SMP Thin node at a glance 257
 - POWER3 SMP Wide node at a glance 257
 - POWER3-II SMP Thin node at a glance 255
 - POWER3-II SMP Wide node at a glance 255
 - pre-install meeting for attached S80/S7A 304
 - publications 246
 - Redundancy 311
 - RS/6000 model equivalents for nodes 310
 - SEPBU 252
 - serial ports 314
 - short frame 248
 - SP Attached Servers
 - Floor Placement Considerations 301
 - the S-Series 301
 - SP Switch (# 4011) 292
 - SP Switch MX Adapter (# 4022) 292
 - SP Switch MX2 Adapter (# 4023) 292
 - SP Switch-8 Adapter (# 4008) 292
 - SP-Attached Servers

- HACWS 302
- Maximum number of SP-Attached Servers 303
- Switchless Attachment 303
- supervisor card 249
- supported AIX and PSSP levels 309
- supported Control Workstations 294
- Switch 286
- Switch Adapter 290
- Switch Board 289
- Switch Chip 288
- Switch Link 286
- Switch MX Adapter (# 4022) 313
- Switch Ports 286
- Switch Router
 - Advantages 300
 - Gigabit Ethernet Media Cards 299
- Switch Routers 298
- Switch system 291
- Switch types 292
- System Partitioning 313
- system partitioning 309
- tall frame 248
- Thin Nodes 275
- user space protocol 309
- uses 245
- Web sites 246
- Wide Nodes 265
- Withdrawn Nodes and their RS/6000 Equivalents 310
- SP 332 MHz SMP Thin node
 - at a glance 258
 - disk requirements and options 284
 - facts and features summary 22
 - mandatory prerequisites 282
 - memory requirements and options 283
 - power system upgrade 281
 - processor requirements and options 283
 - standard features 282
 - supported levels of AIX and PSSP 309
 - switch adapter requirements 284
- SP 332 MHz SMP Wide node
 - adapter placement restrictions 271
 - at a glance 258
 - bus description 271
 - disk requirements and options 274
 - facts and features summary 22
 - mandatory prerequisites 272
 - memory requirements and options 273
 - power system upgrade 271
 - processor requirements and options 273
 - standard features 272
 - supported levels of AIX and PSSP 309
 - switch adapter requirements 274
- SP Attached Servers 301
- SP attachment
 - S70 and S7A 171
- SP Model 500 248, 250
- SP Model 550 248, 250
- SP node Ethernet boot 207
- SP nodes
 - 332 MHz SMP Thin node 281
 - 332 MHz SMP Wide node 270
 - at a glance 255
 - introduction 253
 - performance data 643
 - selection 254
- SP POWER3 SMP Thin node
 - at a glance 257
 - single node configuration rules 285
 - supported levels of AIX and PSSP 309
- SP POWER3 SMP Wide node
 - at a glance 257
 - supported levels of AIX and PSSP 309
- SP POWER3-II SMP Thin node
 - at a glance 255
- SP POWER3-II SMP Wide node
 - at a glance 255
- SP Switch (# 4011) 292
- SP System Attachment Adapter 328
- SP1 244
- SP2 245
- SP332 MHz SMP Thin node
 - bus description 281
- SP-attached server 155
- SPCN
 - Model S7A 165
- SPEC
 - benchmark data for SMP systems 647
 - benchmark data for uni-processor systems 643
 - benchmark definitions 654
 - SPECweb96 benchmark data 650
- SPS switch 206
- SSA
 - see also Communications and I/O Adapters chapter 403
 - architecture 423
- SSA adapters

- 6215 SSA RAID 5 Adapter with 6222 SSA Fast-Write Cache Option 357
- SSA Fast-Write Cache Option (# 6222) 357
- standard features
 - Model 140 28
 - Model 150 35, 61
 - Model 170 41
 - Model 260 47
 - Model 270 53
 - Model F80 85
- standards
 - ANSI X3.131-1986 391
 - X3T9.2/375R Revision 10K 392
- storage adapters 323
 - see adapters, storage
- storage internal
 - disk drives 407
- storage, external
 - architectures
 - Seascope Storage Enterprise Architecture 421
 - Serial Storage Architecture (SSA) 423
 - disk 424
 - 7133-D40 Serial Disk System, Advanced 430
 - 7133-T40 Serial Disk System, Advanced 430
 - 7203 portable disk drive 449
 - positioning 424
 - summary table 425
 - optical devices 480
 - 3995 optical library 482
 - 7210 optical CD-ROM drive 480
 - storage management solutions 486
 - Network Storage Manager 486
 - tape automation products 458
 - 3490E Magnetic Tape Subsystem 461
 - 3494 tape library 466
 - 3570 tape subsystem 471
 - 3575 tape library dataserver 473
 - 3590 tape subsystem 463
 - 7331 8 mm tape library 477
 - 7332 4 mm DDS tape autoloader 475
 - 7337 digital linear tape library 479
 - positioning 458
 - summary table 458
 - tape drives 450
 - 7205 digital linear tape drive 451
 - 7206 external 4 mm DAT tape drive 453
- 7207 1/4" cartridge tape drives 455
- 7208 8 mm tape drive 456
- summary table 450
- superscalar 174
- supervisor card 249
- support processor 294
- supported Control Workstations for the SP 294
- Switch Adapter for the SP 290
- Switch Board for the SP 289
- switch cable 206
- Switch Chip for the SP 288
- switch contention 180
- Switch for the SP 286
- Switch Link for the SP 286
- switch node numbers 206
- Switch Ports for the SP 286
- Switch Routers for the SP 298
- Switch system for the SP 291
- switch traffic 180
- Switch types for the SP 292
- symmetric multi processor 627
- synchronous, SCSI 398
- SysKonnnect SK-NET FDDI-UP SAS adapter 331
- system adapters 325, 326
 - HIPPI 325
 - SP System Attachment Adapter 328
- system backplane 162, 182, 184
- system buses 174
- system console 203, 751
- System Data Repository 206
- system firmware 185
- system monitoring, PSSP 606
- system partitioning 309
- system partitioning, PSSP 606
- System Performance Enhancements 539
- System Power Control Network Cables
 - Model S7A 165
- system serial number 199

T

- tagged command queuing, SCSI 392
- tall frame 206, 248, 250
- tape
 - 4mm 419
 - 8mm 419
- tape drives, external 450
 - 7205 digital linear tape drive 451
 - 7206 external 4 mm DAT tape drives 453

- 7207 1/4" cartridge tape drives 455
- 7208 8 mm tape drive 456
- summary table 450
- tape libraries and tape automation products 458
 - 3490E Magnetic Tape Subsystem 461
 - 3494 tape library 466
 - 3570 tape subsystem 471
 - 3575 tape library dataserer 473
 - 3590 tape subsystem 463
 - 7331 8 mm tape library 477
 - 7332 4 mm DDS tape autoloader 475
 - 7337 digital linear tape library 479
 - positioning 458
 - summary table 458
- tape, internal 419
- TaskGuides 508
- TCP Checksum Offload for Gigabit Ethernet 533
- TCP/IP 530
 - Selective Acknowledgements 535
 - TCP/IP Performance Enhancements 534
 - TCP/IP Resolver 532
 - TCP/IP socket options 535
 - TCP/IP Socks Library 533
 - TCP/IP/ARP 535
- technologies, key new 13
 - 64-bit technology 13
 - NUMA-Q technology 315
 - processor technologies 14
- terminal 203
- thread 534
- threads
 - IP incoming packets 534
 - multiple threads in 64-bit apps 510
 - runnable threads 496
 - threads per process 497
- three I/O drawer configuration 187
- ticket-granting service 293
- timing paths 11, 175
- TME 628
- token ring adapters
 - 4959 4/16 Mbps token-ring adapter 348
- TP Ethernet 207
- TPC-C
 - benchmark data 651
 - benchmark definitions 658
- TPC-D
 - benchmark data 652
 - benchmark definitions 658
- Trailblazer 245

- transfer buffers 180
- twisted pair cables 207
- twisted pair Ethernet 207

U

- UDB 540
- Ultimedia Video Capture Adapter 386
- Ultimedia video capture adapter
 - supported machines 386
- Ultra SCSI Wide vs Ultra SCSI 398
- Ultra SCSI, defined 398
- Unsupported levels of HACMP 597
- upgrade paths
 - Models 140, 150. 260 27
- User Login Performance 511
- uses
 - SP 245
- utility loading for SP frames 251

V

- Virtual Frame Buffer 516
- virtual shared disk, Parallel software 608
- virtual shared disk, PSSP 607
- VisualAge C++
 - RAD 576
- VisualAge C++ V4 576
- VisualAge C++V5 574
- visualization tool, Parallel Environment 613
- Volume Group 497
- VSD, Parallel software 608

W

- WAN adapters
 - 2708 Eicon ISDN DIVA PRO 2.0 S/T Adapter 325
 - 2943 8-port Asynchronous Adapter 332
 - 2944 128-port Asynchronous Adapter 333
 - 2947 ARTIC960Hx 4-Port Selectable Adapter 335
 - 2948 ARTIC960Hx 4-port T1/E1 Adapter 336
 - 2962 2-port Multiprotocol X.25 Adapter 338
 - 6310 ARTIC960RxD Quad Digital Trunk Adapter 337
- Web Browser 493
- Web Server 493
- Web-based 3-D graphics 516
- Web-based System Manager 508

- WebFS 539
- wide SCSI 397
- withdrawn graphics accelerators 387
- WLM 496
- workgroup server 25
- Workgroup Servers (Entry) 25
- working set 10, 175
- Workload Manager 508
- Workstations 25

X

- X.25 2-port Multiprotocol Adapter (# 2962) 338
- X11R5 515
- X11R6 515
 - Broadway release 515
 - new functions 515
- XL Fortran 581
- XL Fortran RTE 580
- XL Fortran V7.1 578
- XL High Performance Fortran 580
- Xmark93, benchmark 368

Y

- yottabyte 10

IBM Redbooks Review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at ibm.com/redbooks
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number	SG24-5120-01
Redbook Title	RS/6000 Systems Handbook 2000 Edition
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	<input type="radio"/> Very Good <input type="radio"/> Good <input type="radio"/> Average <input type="radio"/> Poor
Please identify yourself as belonging to one of the following groups:	<input type="radio"/> Customer <input type="radio"/> Business Partner <input type="radio"/> Solution Developer <input type="radio"/> IBM, Lotus or Tivoli Employee <input type="radio"/> None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	<input type="checkbox"/> Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. ibm.com/privacy/yourprivacy/



Redbooks

RS/6000 Systems Handbook 2000 Edition



RS/6000 Systems Handbook 2000 Edition



Redbooks

The ideal deskside reference for the latest RS/6000 models and features

The RS/6000 System Handbook 2000 edition, SG24-5120, is a comprehensive single-source guide covering the RS/6000 product line. Major hardware and software offerings are introduced and their prominent functions discussed.

Hundreds of tables and figures to accelerate your research

This publication is suitable for professionals wishing to acquire a better understanding of RS/6000 products, including:

- Customers
- Sales and marketing professionals
- Technical support professionals
- IBM Business Partners

A required reading for all RS/6000 and AIX professionals

Inside this publication, you will find:

- A historical look at RS/6000 hardware
- An overview of the latest RS/6000 models
- A short discussion on hardware architecture
- Information on storage, graphics, and communications
- A description of AIX and supported software platforms
- Hundreds of tables and figures

The introduction of this redbook expands the current set of RS/6000 handbooks by providing an ideal, comprehensive, desktop reference that covers the entire product range from the desktop to the raised floor.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-5120-01

ISBN 073841933