



IBM ™ *@server* pSeries Systems Handbook 2003 Edition

The ideal deskside reference for the latest pSeries models and features

Hundreds of tables and figures to accelerate your research

A required reading for all pSeries and AIX professionals



Scott Vetter
Flavio Cappelletti
JaeHoon Chung
Rosario Delgado
Kimberly Fidler
Hansjörg Schneider

ibm.com/redbooks

Redbooks



International Technical Support Organization

**IBM @server pSeries Systems Handbook
2003 Edition**

December 2003

Note: Before using this information and the product it supports, read the information in “Notices” on page xxxi.

Third Edition (December 2003)

This edition applies to IBM @server pSeries and RS/6000 servers as configured and used with AIX 5L, program number 5765-E62.

© Copyright International Business Machines Corporation 1999, 2000, 2003. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	xvii
Tables	xxiii
Notices	xxxix
Trademarksxxxix
Preface	xxxv
The team that wrote this redbook	xxxv
Become a published authorxxxvii
Comments welcomexxxvii
Chapter 1. Introduction to IBM @server pSeries systems	1
1.1 Overview of pSeries systems	2
1.2 Processor descriptions	3
1.2.1 32-bit versus 64-bit computing	3
1.2.2 POWER1	4
1.2.3 POWER2	4
1.2.4 PowerPC	5
1.2.5 RS64	6
1.2.6 POWER3	6
1.2.7 POWER4	7
1.3 Autonomic computing	10
1.3.1 e-business on demand	12
1.3.2 Reliability, Availability, and Serviceability (RAS) features	14
1.3.3 Capacity Upgrade on Demand	16
Chapter 2. Facts and features reference	19
Chapter 3. Entry systems	31
3.1 7028 Model 6C1 and 6E1 IBM @server pSeries 610	32
3.1.1 Product positioning	32
3.1.2 Highlights	33
3.1.3 Technical overview	34
3.1.4 RAS features	36
3.1.5 Minimum and standard features	36
3.1.6 System expansion	38
3.1.7 Features	38
3.1.8 Configuration notes	38

3.1.9	Express configurations	44
3.2	7029 Model 6C3 and 6E3 IBM @server pSeries 615	46
3.2.1	Product positioning	48
3.2.2	Highlights	49
3.2.3	Technical overview	50
3.2.4	RAS features	55
3.2.5	Minimum and standard features	56
3.2.6	System expansion	57
3.2.7	Features	58
3.2.8	Configuration notes	58
3.2.9	Express configurations	65
3.3	7028 Model 6C4 and 6E4 IBM @server pSeries 630	77
3.3.1	Product positioning	78
3.3.2	Highlights	79
3.3.3	Technical overview	80
3.3.4	RAS features	86
3.3.5	Minimum and standard features	87
3.3.6	System expansion	88
3.3.7	Features	89
3.3.8	Configuration notes	89
3.3.9	Express configurations	98
3.4	9114-275 IBM IntelliStation POWER 275	109
3.4.1	General description	109
3.4.2	Configuration notes	110
3.4.3	Express configurations	111
3.5	7043 Model 150 IBM RS/6000 43P	112
3.5.1	Product positioning	113
3.5.2	Highlights	113
3.5.3	Technical overview	114
3.5.4	RAS features	116
3.5.5	Minimum and standard features	116
3.5.6	System expansion	117
3.5.7	Features	118
3.5.8	Configuration notes	118
3.6	7044 Model 170 IBM RS/6000 44P	122
3.6.1	Product positioning	124
3.6.2	Highlights	125
3.6.3	Technical overview	126
3.6.4	RAS features	129
3.6.5	Minimum and standard features	129
3.6.6	System expansion	130
3.6.7	Features	130
3.6.8	Configuration notes	130

3.7	9112-265 IBM IntelliStation POWER 265	134
3.7.1	Product positioning	136
3.7.2	Highlights	136
3.7.3	Technical overview and RAS features	136
3.7.4	Minimum and standard features	136
3.7.5	System expansion	138
3.7.6	Features	138
3.7.7	Configuration notes	138
3.7.8	Express configurations	139
3.8	7026 Model B80 IBM @server pSeries 640	139
3.8.1	Product positioning	140
3.8.2	Highlights	140
3.8.3	Technical overview	141
3.8.4	RAS features	143
3.8.5	Minimum and standard features	143
3.8.6	System expansion	144
3.8.7	Features	145
3.8.8	Configuration notes	145
3.9	Features for entry-level systems	149
Chapter 4. Midrange systems		171
4.1	7038-6M2 IBM @server pSeries 650 server	172
4.1.1	Highlights of the p650	173
4.1.2	Technical overview of the p650	173
4.1.3	RAS features	175
4.1.4	Minimum and standard features p650	176
4.1.5	System expansion of the p650	178
4.1.6	Features of the p650	179
4.1.7	Configuration notes	187
4.1.8	Express configurations	195
Chapter 5. IBM rack-mounted expansion drawers		199
5.1	7311-D20 Expansion Drawer	200
5.1.1	7311-D20 I/O drawer overview	200
5.1.2	Highlights	203
5.1.3	Configuration notes	204
5.1.4	6C4 I/O drawer attachment notes	204
5.1.5	6M2 I/O drawer attachment notes	205
5.2	7311-D10 I/O Drawer	206
5.2.1	Highlights	207
5.2.2	I/O drawer attachment notes	208
5.2.3	Features	209
Chapter 6. High-end servers: IBM @server pSeries 670 and 690		215

6.1	Positioning pSeries 670 and 690	216
6.2	7040 pSeries technical overview	217
6.2.1	Central Electronics Complex (CEC 7040-671/681).	219
6.2.2	Memory subsystem	228
6.2.3	I/O drawer - (7040-61D)	232
6.2.4	Media drawer	236
6.2.5	System Rack - (7040-61R)	237
6.2.6	IBM Hardware Management Console for pSeries.	240
6.2.7	RAS features.	241
6.3	7040-671 IBM @server pSeries 670	242
6.3.1	Highlights of pSeries 670	243
6.3.2	7040 pSeries 670 minimum and standard features	244
6.3.3	7040 pSeries 670 system expansion	244
6.3.4	Configuration notes.	245
6.4	7040-681 IBM @server pSeries 690	266
6.4.1	Highlights pSeries 690	266
6.4.2	7040 pSeries 690 minimum and standard features	267
6.4.3	7040 pSeries 690 system expansion	268
6.4.4	Configuration notes.	269
6.5	Features of the 7040 pSeries 670/690	295
Chapter 7. Racks and rack solution.		309
7.1	Model 7014-T00 overview	310
7.2	Model highlights 7014-T00	311
7.2.1	Description	311
7.2.2	Side panels and side to side rack suite attachment	311
7.2.3	Front door and front trim kit.	312
7.2.4	Ruggedized rack feature	312
7.2.5	Other standard equipment and optional features	312
7.2.6	Physical specifications, operating environment, and power	312
7.2.7	Publications.	317
7.3	Model 7014-T42 overview	317
7.4	Model highlights 7014-T42	318
7.4.1	Description	318
7.4.2	Side panels and side to side rack suite attachment	319
7.4.3	Front door and front trim kit.	319
7.4.4	Ruggedized rack feature	319
7.4.5	Other standard equipment and optional features	320
7.4.6	Physical specifications	320
7.5	Model 7316-TF2 overview.	325
7.5.1	Model highlights 7316-TF2	325
7.5.2	Description	326
7.5.3	Physical specifications	326

7.6	Uninterruptible Power System (UPS)	326
7.6.1	Rack rules	328
7.6.2	Customer setup (CSU)	328
7.6.3	Devices supported	329
7.6.4	Battery runtime	329
7.6.5	UPS Model 9910-Axx	330
7.6.6	UPS Model 9910-Pxx	333
7.6.7	UPS Model 9910-Mxx	338
Chapter 8. High density cluster servers		343
8.1	7039-651 IBM @server pSeries 655 Server	344
8.1.1	The pSeries 655 as a Cluster Server	344
8.1.2	The p655 Clustered System	346
8.1.3	Software requirements	348
8.1.4	p655 Cluster considerations	349
8.1.5	7040-61D I/O Drawer	350
8.1.6	Disk and boot devices	351
8.1.7	7040-W42 System Rack	352
8.1.8	Feature codes on the p655	353
8.2	The RS/6000 SP in the IBM @server Cluster 1600	359
8.2.1	The 9076 555/557 and 556/558 Models	359
8.3	7045-SW4 High-Performance Switch	362
Chapter 9. Storage, communication, and other I/O adapters		365
9.1	PCI storage adapters	368
9.1.1	PCI-X Dual Channel Ultra320 SCSI RAID Adapter (FC 5703)	368
9.1.2	Dual Channel SCSI RAID Enablement Card (FC 5709)	370
9.1.3	PCI-X Dual Channel Ultra320 SCSI Blind Swap (FC 5710)	371
9.1.4	PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap (FC 5711)	372
9.1.5	PCI-X Dual Channel Ultra320 SCSI Adapter (FC 5712)	374
9.1.6	PCI 4-Channel Ultra3 SCSI RAID Adapter (FC 2498)	375
9.1.7	PCI Dual Channel Ultra3 SCSI Adapter (FC 6203)	377
9.1.8	PCI Universal Differential Ultra SCSI Adapter (FC 6204)	377
9.1.9	PCI Dual Channel Ultra2 SCSI Adapter (FC 6205)	378
9.1.10	PCI Single-Ended Ultra SCSI Adapter (FC 6206)	379
9.1.11	PCI Differential Ultra SCSI Adapter (FC 6207)	381
9.1.12	Advanced SerialRAID Plus Adapter (FC 6230)	382
9.1.13	Gigabit Fibre Channel Adapter for 64-bit PCI Bus (FC 6228)	386
9.2	System adapters	386
9.2.1	IBM Short-wave Serial HIPPI PCI Adapter (FC 2732)	386
9.2.2	S/390 ESCON Channel Adapter (FC 2751)	387
9.2.3	IBM RS/6000 SP System Attachment Adapter (FC 8396)	388
9.2.4	SP Switch2 PCI Attachment Adapter (FC 8397)	389

9.2.5	SP Switch2 PCI-X Attachment Adapter (FC 8398)	389
9.3	Asynchronous adapters	390
9.3.1	8-Port Asynchronous Adapter EIA 232/RS-422 (FC 2943)	390
9.3.2	128-Port Asynchronous Adapter EIA-232 (FC 2944)	391
9.4	ARTIC adapters	393
9.4.1	ARTIC960Hx 4-Port Selectable Adapter (FC 2947)	393
9.4.2	ARTIC960RxD Quad Digital Trunk Adapter (FC 6310)	394
9.5	WAN adapters	397
9.5.1	2-Port Multiprotocol X.25 Adapter (FC 2962)	397
9.6	Asynchronous transfer mode adapters	398
9.6.1	Turboways 622 Mbps PCI MMF ATM Adapter (FC 2946)	399
9.6.2	IBM 64-bit/66 MHz PCI ATM 155 UTP Adapter (FC 4953)	399
9.6.3	IBM 64-bit/66 MHz PCI ATM 155 MMF Adapter (FC 4957)	399
9.7	Ethernet and token-ring adapters	400
9.7.1	Gigabit Ethernet - SX Adapter (FC 2969)	400
9.7.2	10/100/1000 Base-T Ethernet PCI Adapter (FC 2975)	401
9.7.3	4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951)	402
9.7.4	4/16 Mbps Token-Ring Adapter (FC 4959)	403
9.7.5	IBM 4-Port 10/100 Ethernet Adapter (FC 4951)	404
9.7.6	IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961)	405
9.7.7	10/100 Mbps Ethernet PCI Adapter II (FC 4962)	406
9.7.8	IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700)	407
9.7.9	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701)	407
9.7.10	2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5706)	408
9.7.11	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707)	409
9.7.12	2 Gigabit Fibre Channel PCI-X Adapter (FC 6239)	410
9.8	Cryptographic adapters	411
9.8.1	IBM e-business Cryptographic Accelerator (FC 4960)	411
9.8.2	PCI Cryptographic Coprocessor-FIPS-4 (FC 4963)	412
Chapter 10.	Graphics accelerators	415
10.1	Available graphics accelerators	417
10.1.1	POWER GXT135P graphics accelerator (FC 2848)	417
10.1.2	POWER GXT4000P graphics accelerator (FC 2826)	419
10.1.3	POWER GXT4500P graphics accelerator (FC 2842)	421
10.1.4	POWER GXT6000P graphics accelerator (FC 2827)	423
10.1.5	POWER GXT6500P graphics accelerator (FC 2843)	425
10.1.6	Display and cable matrix	427
10.2	Withdrawn graphics accelerators	429
Chapter 11.	Internal storage devices	433
11.1	DVD and CD-ROM	434
11.2	Internal tape drives	440

11.3 SCSI drives	444
11.3.1 SCSI-I	444
11.3.2 SCSI-II	444
11.3.3 SCSI-III	445
11.3.4 Overview of SCSI-III standards	445
11.3.5 General terminology	448
11.3.6 Available SCSI disk drives	449
Chapter 12. External storage architectures and devices	451
12.1 IBM external disk storage	452
12.2 Disk systems summary	452
12.2.1 IBM 7133 Serial Disk System Advanced Models T40 and D40	455
12.2.2 Expandable Storage Plus 2104	458
12.2.3 Enterprise storage server 2105 Model 800	461
12.2.4 IBM TotalStorage FASiT900 Storage Server	465
12.2.5 IBM TotalStorage FASiT700 Storage Server	470
12.2.6 IBM TotalStorage FASiT600 Storage Server	473
12.2.7 IBM TotalStorage FASiT200 Storage Server	476
12.3 IBM TotalStorage Tape drive products	479
12.3.1 7205 IBM External Digital Linear Tape Drive	480
12.3.2 7206 IBM External 4 mm Tape Drive	481
12.3.3 7207-122 IBM 4 GB external SLR5 QIC tape drive Model 122	484
12.4 IBM Tape Libraries/Autoloaders	485
12.4.1 3590 IBM TotalStorage Enterprise Tape System	485
12.4.2 3494 IBM TotalStorage Enterprise Automated Tape Library	487
12.5 LTO Ultrium tape drives, libraries, and autoloaders	491
12.5.1 LTO product positioning notes	491
12.5.2 3580 IBM Ultrium Tape Drive	492
12.5.3 3581 IBM Ultrium Tape Autoloader	496
12.5.4 3582 IBM TotalStorage Ultrium Tape Library Model L23	498
12.5.5 3583 IBM Ultrium Scalable Tape Library	501
12.5.6 3584 IBM UltraScalable Tape Library	507
12.6 The IBM SAN solution	511
12.6.1 IBM @server pSeries Fibre Channel Host Bus Adapters (HBA)	512
12.6.2 IBM TotalStorage SAN switches	513
12.6.3 Sample configurations	517
Chapter 13. Operating systems	521
13.1 AIX 5L Version 5.1	522
13.1.1 Hardware requirements	522
13.1.2 Software requirements	524
13.1.3 AIX 5L Version 5.1 release and support date	525
13.1.4 AIX 5L Version 5.1 highlights	525

13.1.5 Scalability and system management	525
13.1.6 Resource management	535
13.1.7 Storage management	537
13.1.8 Reliability, Availability, and Serviceability (RAS)	539
13.1.9 Development and performance tools	542
13.1.10 Linux affinity	545
13.1.11 Network	546
13.1.12 AIX system security	550
13.1.13 Networking security	554
13.2 AIX 5L Version 5.2	555
13.2.1 Hardware requirements	555
13.2.2 Software requirement	561
13.2.3 AIX 5L Version 5.2 release and support date	563
13.2.4 AIX 5L Version 5.2 highlights	563
13.2.5 System management	563
13.2.6 Resource management	565
13.2.7 Storage management	567
13.2.8 Reliability, Availability, and Serviceability (RAS)	569
13.2.9 Development and performance tools	570
13.2.10 Linux affinity	571
13.2.11 Network	571
13.2.12 AIX system security	573
13.2.13 Networking security	574
13.3 AIX 5L Expansion, Bonus, and Web Download Pack	575
13.3.1 System management	576
13.3.2 e-business and e-commerce	579
13.3.3 Development	580
13.3.4 Multimedia	582
13.3.5 Interoperability	582
13.3.6 Storage management	583
13.3.7 Contents summary	584
13.3.8 How to order	588
13.4 Linux	589
13.4.1 Introduction to Linux	589
13.4.2 Linux and pSeries	595
13.4.3 Linux affinity on AIX 5L	602
13.4.4 Software service and support	607
Chapter 14. Cluster, resource, and performance management software	609
14.1 IBM Cluster Systems Management for AIX (CSM)	610
14.1.1 Product positioning	610
14.1.2 CSM for AIX 5L tasks	611
14.1.3 CSM for AIX Version 1.3	612

14.1.4	Hardware requirements	613
14.1.5	Software requirements	614
14.2	IBM Parallel System Support Programs (PSSP)	615
14.2.1	Product positioning	615
14.2.2	PSSP Version 3.5 functional enhancements	616
14.2.3	Hardware requirements	617
14.2.4	Software requirements	617
14.3	General Parallel File System for AIX 5L	618
14.3.1	Product positioning	618
14.3.2	Hardware requirements	620
14.3.3	Software requirements	621
14.4	AIX Workload Manager	621
14.5	AIX Performance ToolBox and Performance AIDE	624
14.5.1	Hardware requirements	625
14.5.2	Software requirements	625
14.6	AIX Fast Connect for POWER Version 3.1	625
Chapter 15. Cluster systems		627
15.1	HACMP	628
15.1.1	Basic concepts of HACMP	628
15.1.2	Components	629
15.1.3	Resource groups	632
15.1.4	Cluster configurations	634
15.1.5	Features in HACMP	638
15.1.6	HACMP history	640
15.1.7	HACMP and HACMP ES Version 4.4	647
15.1.8	HACMP and HACMP ES Version 4.4.1	648
15.1.9	HACMP and HACMP ES Version 4.5	649
15.1.10	HACMP Version 5.1	650
15.1.11	Hardware supports matrix	652
15.1.12	Software requirements	660
15.2	HAGEO/GeoRM	663
15.2.1	Overview	664
15.2.2	Hardware requirements	665
15.2.3	Software requirements	667
15.2.4	Configuration examples	667
15.2.5	New features of HAGEO Version 2.4/GeoRM Version 2.4	670
15.3	Cluster 1600	672
15.3.1	Overview	672
15.3.2	Parallel System Support Program Version 3.5	673
15.3.3	Cluster Systems Management Version 1.3.1	674
15.3.4	General Parallel File System for AIX Version 2.1	674
15.3.5	Models and features	674

15.3.6 Cluster enhancements	676
15.3.7 Cluster 1600 with PSSP	678
15.3.8 Cluster 1600 with CSM	681
Chapter 16. Resource management	685
16.1 Partitioning	686
16.1.1 Logical partitioning overview	686
16.1.2 Partitioning implementations	686
16.1.3 Partitioning support on pSeries servers	688
16.1.4 LPAR configuration notes	690
16.1.5 Terminology used in partitioning	695
16.1.6 Four terms regarding memory	697
16.2 Partitioning implementation on pSeries servers	701
16.2.1 Partitioning implementation	701
16.2.2 Partition resources	712
16.2.3 I/O device assignment considerations	717
16.2.4 Service authority	724
16.3 Dynamic logical partitioning	725
16.3.1 Dynamic logical partitioning overview	725
16.4 Managing partitions	727
16.4.1 Hardware Management Console (HMC) user interface	728
16.4.2 HMC graphical user interface	728
16.5 IBM Hardware Management Console for pSeries	729
16.6 Electronic Service Agent	735
16.6.1 Overview	735
16.6.2 Functions	738
16.6.3 System support	739
16.7 Service Focal Point	739
16.7.1 Overview	739
16.7.2 Using the Hardware Service Function	741
16.8 Inventory Scout Services	743
16.8.1 Overview	743
16.8.2 Functions	743
Appendix A. Site and hardware planning information	747
7026 Model B80	748
7028 Model 6C1 and 6E1	750
7028 Model 6C4 and 6E4	751
7029 Model 6C3 and 6E3	753
7038-6M2 p650	754
7039 pSeries 655	756
7040 pSeries 670	762
7040 pSeries 690	768

7043 43P Model 150	773
7044 44P Model 170	775
7311 Model D10	777
7311 Model D20	778
9112 Model 265	779
9114 Model 275	780
7014 Model T00 and T42 Rack	782
Appendix B. Adapter placement guidelines	785
Introduction	787
System performance	787
Integrated adapters	787
32-bit versus 64-bit PCI slots	787
33 MHz versus 50/66 MHz 64-bit PCI slots	787
Connectivity versus performance overview	788
Other restrictions	789
7025 pSeries 620 Models 6F0 and 6F1	789
7026 pSeries 640 Model B80	798
7028 pSeries 610 Models 6C1 and 6E1	803
7028 pSeries 630 Models 6C4 and 6E4 (4-slot PCI riser)	809
Logical partition (LPAR) considerations	809
7028 pSeries 630 Models 6C4 and 6E4 (6-Slot PCI Riser)	815
7029 pSeries 615 Models 6C3 and 6E3	820
Logical partition (LPAR) considerations	820
7040 pSeries 670 Model 671 drawers 61D	824
Logical partition (LPAR) considerations	827
7040 pSeries 690 Model 681 drawer 61D	835
Logical partition (LPAR) considerations	837
7038 pSeries 650	845
Logical partition (LPAR) considerations	846
7039 pSeries 655 Model 651	850
Logical partition (LPAR) considerations	850
7043 Model 150	859
7043 Model 150 multiple graphics adapter placement guide	864
7044 Model 170	867
7044 Model 170 multiple graphics adapter placement guide	874
7044 Model 270	876
7044 Model 270 multiple graphics adapter placement guide	882
7046 Model B50	884
7311 Model D10	887
7311 Model D20	891
9076 RS/6000 SP systems	897
9112-265 IntelliStation POWER 265	897

9114-275 IntelliStation POWER 275	902
High-performance adapters	904
Appendix C. Power cord features	907
Appendix D. SCSI cabling examples	917
Cabling the PCI single-ended Ultra SCSI adapter	918
Supported Ultra configurations for this adapter	919
Cables and terminators for single-ended SCSI adapters	919
Single-Ended adapter-to-first device cables	920
Device-to-device cables	920
System-to-system cables	921
Terminators for use with this adapter	922
Cabling examples for the PCI Ultra SCSI adapter	922
Cabling the PCI Differential-Ended Ultra SCSI adapter	924
Cables and terminators for Single-Ended SCSI adapters	924
Single-ended adapter-to-first device cables	924
SCSI Differential cable lengths using this adapter	925
Differential-ended adapter-to-first device cables	925
Device-to-device cables	926
Terminators for use with this adapter	927
Examples for the PCI Differential Ultra SCSI Adapter	928
Appendix E. Supported peripherals by device matrix	931
Appendix F. Customer installation matrix and processor groups	947
Appendix G. pSeries performance	949
Performance of pSeries systems	950
Performance benchmarks	950
System Performance Evaluation Corporation (SPEC)	951
Transaction oriented benchmarks	952
Relative Performance (rPerf)	952
NotesBench benchmark	953
Appendix H. System life time information	955
7006 IBM RS/6000 server	956
7007 IBM RS/6000 server	956
7008 IBM RS/6000 server	956
7009 IBM RS/6000 compact server	957
7010 IBM RS/6000 server	957
7011 IBM RS/6000 server	957
7012 IBM RS/6000 server	958
7013 IBM RS/6000 server	959

7014 IBM RS/6000 server	960
7015 IBM RS/6000 server	960
7017 IBM RS/6000 enterprise server	961
7024 IBM RS/6000 server	962
7025 IBM RS/6000 server	962
7026 IBM RS/6000 server	963
7027 IBM high capacity drawer	963
7028 IBM @server pSeries 610 server	964
7030 IBM RS/6000 server	964
7038-6M2 IBM @server pSeries 650	964
7039-651 IBM @server pSeries 655 Server Model 651	965
7040 IBM @server pSeries 670/690 servers	965
7043 IBM RS/6000 43P server	965
7044 IBM RS/6000 server	966
7046-B50 IBM RS/6000 server	966
7202-900 IBM RS/6000 expansion rack Model 900	966
7236-001 IBM RS/6000 Media Streamer Model 001	967
7248 IBM RS/6000 43P series	967
7311 IBM Rack-mounted Drawers	967
7315-C01 IBM Hardware Management Console	968
7316 IBM Rack-mounted Console	968
7317 IBM RS/6000 Telecom Server and External SCSI Disk	968
7318 IBM Serial Communications Network Server	969
7319 IBM RS/6000 Fibre Channel Switch and Adapter	969
Appendix I. Firmware	971
Abbreviations and acronyms	973
Related publications	983
IBM Redbooks	983
AIX product manuals	983
Other publications	984
Online resources	984
How to get IBM Redbooks	991
Help from IBM	992
Index	993

Figures

1-1	pSeries product line	2
1-2	pSeries microprocessor history	3
1-3	POWER4 chip logical view	9
1-4	Overview of an on demand operating environment	14
2-1	pSeries and IntelliStation products in focus	19
3-1	p610 Model 6C1 and 6E1	32
3-2	p610 Model 6C1 and 6E1 general view	33
3-3	p610 Model 6C1 and 6E1 system block diagram	35
3-4	Memory card locations	39
3-5	p610 Model 6E1 disks and media bay locations	41
3-6	p610 Model 6C1 disks and media bay locations	42
3-7	p615 Model 6E3 and 6C3 physical packaging	48
3-8	p615 Model 6C3 rack view and Model 6E3 deskside view	49
3-9	Conceptual diagram of the p615 POWER4+ system architecture	51
3-10	The p615 processor/cache module	52
3-11	Enlarged view of the processor subsystem on system planar	52
3-12	Conceptual diagram of POWER4+ processor and memory subsystem	53
3-13	Main storage ECC and extensions	54
3-14	Models 6C4 and 6E4 (POWER4+ system with six PCI-X slots)	59
3-15	Memory placement for the p615	60
3-16	IBM barcode label	61
3-17	p615 bay locations	63
3-18	p615 Model 6C3 PCI-X slot locations	64
3-19	p630 Model 6C4 and 6E4	78
3-20	p630 Model 6C4 rack view, and Model 6E4 workstation view	79
3-21	Diagram of Models 6C4 and 6E4 POWER4+ system architecture	81
3-22	Comparison between SCM and MCM	82
3-23	POWER4+ processor card layout	82
3-24	Conceptual diagram of POWER4+ processor and memory subsystem	83
3-25	Views of Models 6C4 and 6E4 (POWER4+ system with 6 PCI-X slots)	90
3-26	Memory DIMM slot and LED locations	93
3-27	p630 Model 6C4 and 6E4 bay locations	96
3-28	IntelliStation POWER 275	109
3-29	Model 7043-150 with peripherals	113
3-30	RS/6000 Model 150 diagram	115
3-31	System board locations of RS/6000 Model 150	119
3-32	RS/6000 Model 150 front view	120
3-33	PCI riser card for RS/6000 Model 150	122

3-34	Model 7044-170 with peripherals	124
3-35	System logic flow diagram for 44P Model 170	127
3-36	Memory card locations	131
3-37	Disks and media bay locations	133
3-38	IBM IntelliStation 9112 Model 265	135
3-39	pSeries 640 Model B80 entry rack server	139
3-40	System architecture diagram for p640 Model B80	142
3-41	Memory card locations	146
3-42	Disks and media bay locations	148
4-1	7038-6M2 p650 with one I/O drawer	172
4-2	p650 system architecture	174
5-1	7311-D20 I/O drawer	200
5-2	7311 Model D20 front and rear view	200
5-3	Inside view of 7311 Model D20 I/O drawer	201
5-4	Conceptual diagram of the I/O drawer	202
5-5	Remote I/O and SPCN cabling	203
5-6	7311-D10 I/O drawer enclosure	207
6-1	7040 modular design	217
6-2	CEC front view	220
6-3	CEC rear view	221
6-4	CEC backplane orthogonal view	222
6-5	POWER4 multichip module	226
6-6	Multichip module with L2, L3, and memory	227
6-7	pSeries 690 1.3 GHz Turbo and 1.3 GHz HPC feature	228
6-8	MCM, L3 cache, and memory slots relationship on backplane	229
6-9	MCM, L3 cache, and memory slots relationship on backplane	230
6-10	Primary and secondary I/O books	231
6-11	I/O drawer rear view	233
6-12	Difference between RIO and RIO-2 connectors	234
6-13	Logical view of an RIO drawer	234
6-14	Logical view of an RIO-2 drawer	235
6-15	Media drawer power and SCSI connection	237
6-16	Power subsystem locations in BPA	238
6-17	Graphical user interface on the HMC	241
6-18	Front view pSeries 670	242
6-19	Rack with expansion frame pSeries 690	266
7-1	7014-T00 with pSeries 610 servers	310
7-2	7014-T00 rack with AC power distribution bus type 6	313
7-3	Type 7 power distribution bus	314
7-4	7014-T42 rack	318
7-5	Rear view of a 7014-T42 rack with a type 6 power distribution	321
7-6	Type 7 power distribution bus	322
7-7	7316-TF2 flat panel console	325

8-1	8-way 1.5 GHz and 4-way 1.7 GHz MCMs	345
8-2	Clustered system with p655 servers, Control Workstation, and HMC	347
8-3	p655 in a PSSP Version 3.5 managed cluster	349
10-1	POWER GXT135P	417
10-2	POWER GXT4000P	419
10-3	POWER GXT4500P	421
10-4	POWER GXT6500P	426
12-1	IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer	456
12-2	IBM Expandable Storage Plus SCSI	458
12-3	The Enterprise Storage Server 2105 Model 800	462
12-4	The FAST900 Model 90U	466
12-5	The FAST700 Model 1RU	470
12-6	The FAST600 Model 60U	473
12-7	The FAST200 Model 1RU	476
12-8	7205 External Digital Linear Tape Drive	480
12-9	7206 External 4 mm Tape Drive	481
12-10	7207-122 IBM 4 GB External SLR5 QIC Tape Drive	484
12-11	IBM 3580 Ultrium Tape Drive	492
12-12	IBM 3581Ultrium Tape Autoloader	496
12-13	IBM TotalStorage Ultrium Tape Library 3582 Model L23	498
12-14	3583 Ultrium Scalable Tape Library	501
12-15	3584 UltraScalable Tape Library	507
12-16	IBM SAN Fibre Channel Switch 3534-F08	514
12-17	IBM SAN Fibre Channel Switch 2109-F16	515
12-18	IBM SAN Fibre Channel Switch 2109-F32	516
12-19	2109-M12	517
12-20	Schematic small SAN solution	518
12-21	Schematic medium SAN solution	519
12-22	Schematic large SAN solution	520
13-1	Example of Web-based System Management	528
13-2	AIX and Linux with pSeries	591
13-3	Configuration example	598
13-4	Linux partitions on the pSeries system	601
13-5	AIX toolbox for Linux applications	605
15-1	An example of high availability cluster	631
15-2	Hot-standby configuration	634
15-3	Mutual takeover configuration	636
15-4	Third-party takeover configuration	637
15-5	Two nodes at each site	668
15-6	Two nodes at the primary site and one node at the backup site	669
15-7	One node at each site	669
16-1	Virtual and physical memory relationship	698
16-2	Non-contiguous mapping	700

16-3	Interaction of AIX and firmware in a partition	709
16-4	Partitions, partition profiles, and system profiles	713
16-5	HMC graphical user interface	729
16-6	Picture of a deskside HMC (a rack-mount version is also available) . .	730
16-7	Communication between the HMC and the service processor	732
16-8	Typical Service Agent configuration.	736
16-9	Service Agent with the HMC	737
16-10	Service Agent on the HMC	738
16-11	Error reporting and consolidation.	740
16-12	Service Focal Point: Hardware Service Functions.	741
16-13	Hardware Service Functions overview.	742
16-14	FRU LED management	742
16-15	Inventory Scout Services	744
A-1	Proposed floor layout for multiple 7039 systems	759
A-2	Service clearance for 7039 systems with thin doors	760
A-3	Service clearance for 7039 systems with acoustical doors	761
A-4	Proposed floor layout for multiple systems	765
A-5	Service clearance for 7040 single-frame systems with slimline doors. .	766
A-6	Service clearance for 7040 single-frame systems w/ acoustical doors .	767
A-7	Proposed floor layout for multiple systems	771
A-8	Service clearance for 7040 single/double frame with slimline doors . .	772
A-9	Clearance for 7040 single and double frame with acoustical doors. . .	773
A-10	7043 43P Model 150 physical planning diagram	775
A-11	7044 44P Model 170 physical planning diagram	776
B-1	System 7025 6F0 and 6F1 unit rear view with numbered slots	790
B-2	System 7026 B80 unit rear view with numbered slots	798
B-3	System 7028 6C1 and 6E1 unit rear view with numbered slots.	803
B-4	7028 6C4 and 6E4 rear view (4-slot PCI riser) with numbered slots . .	810
B-5	7028 6C4 and 6E4 rear view (6-slot PCI riser) with numbered slots . .	815
B-6	System 7029 6C3 and 6E3 unit rear view with numbered slots.	821
B-7	System 7040 671 drawer 61D unit rear view with numbered slots . . .	825
B-8	System 7040 681 drawer 61D unit rear view with numbered slots . . .	835
B-9	System 7038 6M2 unit rear view with numbered slots.	846
B-10	System 7039 651 unit rear view with numbered slots	851
B-11	System 7039 651 drawer 61D unit rear view with numbered slots . . .	852
B-12	Model 150 system unit rear view with numbered slots	860
B-13	Model 170 system unit rear view with numbered slots	868
B-14	Model 270 system unit rear view with numbered slots	877
B-15	Model B50 system unit rear view with numbered slots	884
B-16	System 7311 D10 unit rear view with numbered slots.	888
B-17	System 7311 D20 unit rear view with numbered slots.	893
B-18	System 9112 Model 265 unit rear view with numbered slots.	898
B-19	System 9114 Model 275 unit rear view with numbered slots.	902

D-1	Single-ended internal devices cabling	922
D-2	External devices cabling for both narrow and wide bus	923
D-3	Differential External Narrow Bus	928
D-4	Differential External Wide Bus	929
G-1	Example of performance report	951

Tables

2-1	Facts and features for RS/6000 Model 170, and IntelliStation 265, 275	20
2-2	Facts and features for pSeries Model 6C1, 6E1, 6C3, 6E3, and 6E4	21
2-3	Facts and features for pSeries Models 6C4, B80, and 6M2	22
2-4	Facts and features for pSeries Models 655, 760, and 6901	23
2-5	System unit details	25
2-6	System unit details	26
2-7	Server I/O attachment	27
2-8	Peak bandwidth	28
2-9	Standard warranty in United States (other countries may vary)	28
2-10	Standard warranty in United States (other countries may vary)	28
2-11	System software: Model 150 through Model 640	29
2-12	System software: Model 650 through Model 265	29
3-1	p610 Model 6C1 and 6E1 minimum and standard features	37
3-2	p610 Model 6C1 and 6E1 system expansion features	38
3-3	Disk and media bay locations and SCSI-IDs	42
3-4	Internal storage devices supported in systems 6C1 and 6E1	43
3-5	p610 express configuration summary	44
3-6	p615 Model 6C3 and 6C4 minimum and standard features	56
3-7	p615 Model 6C4 and 6E4 system expansion features	58
3-8	Disk and media bay locations and SCSI IDs	63
3-9	p615 express configuration summary	65
3-10	2-way configurations comparative table	83
3-11	GX bus and fabric bus characteristics	85
3-12	p630 Model 6C4 and 6E4 minimum and standard features	87
3-13	p630 Model 6C4 and 6E4 system expansion features	88
3-14	Examples of memory DIMM placement	92
3-15	Example of memory DIMM placement for maximum performance	92
3-16	Disk and media bay locations and SCSI IDs	97
3-17	Graphic adapters supported on systems with 4 and 6 slots	98
3-18	p630 express configuration summary	99
3-19	RS/6000 Model 150 minimum and standard features	116
3-20	Model 150 system expansion	117
3-21	Disk and media bay locations and SCSI-IDs	120
3-22	Model 170 minimum and standard features	129
3-23	Model 170 system expansion	130
3-24	Disk and media bay locations and SCSI-IDs	133
3-25	IntelliStation Model 265 minimum and standard features	137
3-26	IntelliStation Model 265 system expansion features	138

3-27	Maximum number of graphic adapters supported on systems	139
3-28	Model B80 minimum and standard features	143
3-29	Model B80 system expansion features	144
3-30	Disk and media bay locations and SCSI-IDs	148
3-31	Features for entry-level systems	150
4-1	p650 minimum and standard features	176
4-2	Possible maximum processor, memory, and storage features	178
4-3	Features of the p650	179
4-4	Recommended memory configurations	188
4-5	Max. number of adapters supported in combined system 6M2	192
4-6	Adapters that cannot be installed in slot 7	193
4-7	p650 express configuration summary	195
5-1	7311-D20 I/O drawer at a glance	203
5-2	7311-D10 I/O drawer at a glance	207
5-3	Features of the 7311-D20 and D10 IO drawers	209
6-1	Differences between pSeries 670 and pSeries 690	218
6-2	Supported combinations of processors	223
6-3	pSeries 670 minimum and standard feature	244
6-4	pSeries 670 system expansion	244
6-5	Supported memory configurations for pSeries 670	248
6-6	Number of available I/O loops for each combination of MCM	253
6-7	Bulk power regulators required with POWER4 1.1 or 1.3 GHz	256
6-8	Bulk power regulators required for svrs w/ POWER4+ 1.5 GHz MCM	256
6-9	Supported adapters per LPAR within a p670	258
6-10	The pSeries 670 minimum and standard feature	267
6-11	Processor, memory, and storage configurations	268
6-12	Supported memory cards configurations	273
6-13	Available I/O loops	281
6-14	Bulk power regulators with POWER4 1.1 or 1.3GHz	285
6-15	Bulk power regulators with POWER4+ 1.5 or 1.7GHz	285
6-16	Adapter supported per LPAR within a p690	288
6-17	Feature code description pSeries 670/690	296
7-1	Allowable PDU combinations within a single T00 or T42 rack	316
7-2	Allowable PDU combinations within a single T00 or T42 rack	324
7-3	Overview about the several 9910 models	327
7-4	APC application matrix	330
7-5	Help to choose the right Powerware UPS	334
7-6	Choose the right MGE UPS	339
8-1	Recommended p655 memory configurations	345
8-2	Features used in the p655	353
9-1	Available PCI I/O adapters and their feature codes	365
9-2	Overview of the compatibility and use of both Optical Extenders	385
9-3	The 1.2 Mbps RANs and cables	392

9-4	The 2.4 Mbps RANs and cables	392
9-5	ARTIC960Hx cables	394
9-6	ARTIC960 cables and features	395
9-7	Cable information for the 2-Port Multiprotocol PCI Adapter	398
10-1	Features of POWER GXT135P	417
10-2	POWER GXT135P supported systems	418
10-3	Features of POWER GXT4000P	419
10-4	POWER GXT4000P supported systems	420
10-5	Features of POWER GXT4500P	422
10-6	POWER GXT4500P supported systems	423
10-7	Features of POWER GXT6000P	424
10-8	POWER GXT6000P supported systems	425
10-9	Features of POWER GXT6500P	426
10-10	POWER GXT6500P supported systems	427
10-11	Cable feature to graphics accelerator	428
10-12	Withdrawn graphics accelerator	429
11-1	Breakdown of the SCSI-III standards	445
11-2	Available SCSI disk drives	449
12-1	Features of the ESS 2105, 7133 D40/T40, and 2104 DU3/TU3T	452
12-2	The key features of the IBM FAST Storage Servers	453
12-3	Rack-mountable D40/Deskside T40 minimum to maximum capacity	457
12-4	The physical dimensions of the IBM TotalStorage 2104 EXP Plus	460
12-5	IBM 7205 External SDLT320 Tape Drive at a glance	481
12-6	IBM 7206-220 Model DDS-4 at a glance	482
12-7	IBM 7206-VX2 Model VX2 at a glance	483
12-8	IBM 7207-122 4 GB SLR5 QIX Tape drive at a glance	484
12-9	3590 tape drive models at a glance	486
12-10	IBM TotalStorage 3494 Tape Library at a glance	490
12-11	IBM Ultrium Tape Drive at a glance using Ultrium 1 drives	493
12-12	3580 IBM Ultrium Tape Drive at a glance using Ultrium 2 drives	494
12-13	IBM 3581 Ultrium Tape Drive at a glance	497
12-14	IBM TotalStorage Ultrium Tape Library 3582 Model L23 at a glance	499
12-15	IBM TotalStorage Ultrium Scalable Tape Library 3583 at a glance	504
12-16	IBM TotalStorage UltraScalable Tape Library 3584 at a glance	509
13-1	Memory and disk requirements	524
13-2	Release and withdrawal date	525
13-3	Locale support	530
13-4	Euro currency symbol	531
13-5	New locales in AIX 5L Version 5.1	532
13-6	Comparison between JFS2 and JFS	539
13-7	Actions at system hang condition	540
13-8	Version 5.2 withdrawn PCI adapter support	556
13-9	Version 5.2 withdrawn PReP-specific ISA adapter support	557

13-10	Version 5.2 withdrawn ISA adapter support	558
13-11	Version 5.2 PCI RS/6000 withdrawn support listing	558
13-12	Version 5.2 MCA RS/6000 withdrawn support listing	559
13-13	Version 5.2 MCA-based SP nodes withdrawn support	560
13-14	Version 5.2 device support withdrawn	561
13-15	Software requirement.	561
13-16	Release and withdrawal date	563
13-17	AIX 5L Expansion Pack	584
13-18	AIX 5L Bonus Pack	585
13-19	AIX 5L Web Download Pack	587
13-20	Product feature codes	589
13-21	Linux distribution support for pSeries 64-bit systems	596
14-1	CSM for AIX 5L Version 1.3.1 service required	614
15-1	Features and benefits of HACMP	639
15-2	HACMP life cycle information	640
15-3	HACMP supported systems	652
15-4	Communication adapters	655
15-5	External storage subsystems	658
15-6	Router support	660
15-7	Rack-mounted storage subsystems	660
15-8	HAGEO releases and features	664
15-9	Features of Cluster 1600	675
15-10	Description of feature codes	675
15-11	Support availability	678
15-12	Cluster limits for PSSP	679
15-13	Cluster limits for CSM	682
16-1	Supported partitioning-capable pSeries servers	688
16-2	Maximum number of processors, memory size, and partitions	689
16-3	Physical memory size and number of allocatable partitions	716
16-4	The minimum firmware level for dynamic logical partitioning	727
A-1	7026 Model B80 site and hardware planning information	748
A-2	Maximum operating temperatures	749
A-3	7028 Model 6C1 and 6E1 site and hardware planning information	750
A-4	7028 Model 6C4 and 6E4 site hardware planning information	751
A-5	7029 Model 6C3 and 6E3 site and hardware planning information	753
A-6	Operating temperature for FC 6134 at various altitudes	754
A-7	7038-6M2 site and hardware planning information	754
A-8	7039 pSeries 655 multiple components	756
A-9	7039 pSeries 655 site and hardware planning information	756
A-10	7040 pSeries 670 multiple components	762
A-11	7039 pSeries 670 site and hardware planning information	762
A-12	7040 pSeries 690 multiple components	768
A-13	7039 pSeries 690 site and hardware planning information	769

A-14	7043 43P Model 150 site and hardware planning information	773
A-15	7044 44P Model 170 site and hardware planning information	775
A-16	7311 Model D10 site and hardware planning information	777
A-17	7311 Model D20 site and hardware planning information	778
A-18	9112 Model 265 site and hardware planning information	779
A-19	9114 Model 275 site and hardware planning information	780
A-20	Model T00 rack site and hardware planning information	782
A-21	Model T42 rack site and hardware planning information	783
B-1	Placement guidelines 7025 6F0 or 6F1	790
B-2	Placement guidelines 7026 B80	799
B-3	Placement guidelines 7028 6C1 and 6E1	804
B-4	Slot Location Reference 7028 6C4 and 6E4 (4-Slot PCI riser card) . .	810
B-5	Placement guidelines 7028 6C4 and 6E4 (4-slot PCI riser card)) . . .	811
B-6	Slot Location Reference 7028 6C4 and 6E4 (6-Slot PCI riser card) . .	815
B-7	Placement guidelines 7028 6C4 and 6E4 (6-Slot PCI riser card) . . .	816
B-8	Slot Location Reference 7029 6C3 and 6E3	821
B-9	Placement guidelines 7029 6C3 and 6E3	822
B-10	Slot Location Reference 7040 pSeries 670 drawer 61D	825
B-11	Placement guidelines 7040 671 drawer 61D	827
B-12	Slot Location Reference for the 7040 681 drawer 61D	836
B-13	Placement guidelines 7040 pSeries 690 drawer 61D	838
B-14	Slot location reference 7038 6M2	847
B-15	Placement guidelines 7038-6M2	847
B-16	Slot location reference 7039 651	851
B-17	Slot Location Reference drawer 61D	852
B-18	Placement guidelines 7039 651	853
B-19	Placement guidelines 7043 150	860
B-20	7043 Model 150 multiple graphics adapter placement table	865
B-21	Placement guidelines 7044 170	868
B-22	7044 Model 170 multiple graphics adapter placement table	874
B-23	Placement guidelines 7044 270	877
B-24	7044 Model 270 multiple graphics adapter placement table	882
B-25	Placement guidelines 7046 B50	884
B-26	Slot location reference 7311 D10	888
B-27	Placement guidelines 7311 D10	889
B-28	Slot location reference 7311 D20	893
B-29	Placement guidelines 7311 D20	894
B-30	Placement guidelines 9112 Model 265	898
B-31	Placement guidelines 9114 Model 275	903
C-1	Power cords list by feature code	908
C-2	Power cord list by feature code	909
C-3	Power cord connector spotters reference	911
D-1	Single-Ended Ultra SCSI adapter-to-first device cables	920

D-2	Device-to-device cables for single-ended installations	920
D-3	16 Bit SCSI-II system-to-system cable	921
D-4	Terminators for single-ended installations	922
D-5	Internal system cables	922
D-6	Cable and terminators for single-ended adapters	923
D-7	Single-Ended Ultra SCSI adapter-to-first device cables	925
D-8	Differential-ended Ultra SCSI adapter-to-first device cables	925
D-9	Device-to-device cables for differential-ended installations	926
D-10	Terminator for differential-ended Installations	928
D-11	Cables and terminators for the PCI Differential Ultra SCSI Adapter . .	929
E-1	External devices for B80, 6C1, 6E1, 6C4, 6E4, 6C3, and 6E3	932
E-2	External devices for 6M2, 651, 671, 681, 150, 170, 265, and 275 . . .	938
F-1	RS/6000 models and Feature Codes set-up information.	948
H-1	7006 IBM RS/6000 server	956
H-2	7007-N40 IBM RS/6000 Model N40 server	956
H-3	7008 IBM RS/6000 server	956
H-4	7009 IBM RS/6000 compact server	957
H-5	7010 IBM RS/6000 server	957
H-6	7011 IBM RS/6000 server	957
H-7	7012 IBM RS/6000 server	958
H-8	7013 IBM RS/6000 server	959
H-9	7014 IBM RS/6000 server	960
H-10	7015 IBM RS/6000 server	960
H-11	7017 IBM RS/6000 enterprise server.	961
H-12	7024 IBM RS/6000 server	962
H-13	7025 IBM RS/6000 server	962
H-14	7026 IBM RS/6000 server	963
H-15	7027 IBM high capacity drawer	963
H-16	7028 IBM @server pSeries 610 server	964
H-17	7030 IBM RS/6000 server	964
H-18	7038-6M2 IBM @server pSeries 650 server	964
H-19	7039-651 IBM @server pSeries 655 Model 651	965
H-20	7040 IBM @server pSeries 670/690 servers	965
H-21	7043 IBM RS/6000 43P server	965
H-22	7044 IBM RS/6000 server	966
H-23	7046-B50 IBM RS/6000 Model B50.	966
H-24	7202-900 IBM RS/6000 expansion rack Model 900	966
H-25	7236-001 IBM RS/6000 Media Streamer Model 001	967
H-26	7248 IBM RS/6000 43P series.	967
H-27	7311 IBM Rack-Mounted Drawers.	967
H-28	7315-C01 IBM Hardware Management Console C01	968
H-29	7316 IBM Rack-Mounted Console	968
H-30	7317 IBM RS/6000 Telecom Server and External SCSI Disk	968

H-31	7318 IBM Serial Communications Network Server	969
H-32	7319 IBM RS/6000 Fibre Channel Switch and Adapter	969

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include 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.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in

any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®	MVS/ESA™	QMF™
@server®	MVS™	Redbooks (logo)  ™
AFS®	Netfinity®	Redbooks™
AIX 5L™	NetView®	RISC System/6000®
AIX®	Notes®	RS/6000®
AS/400®	OS/2®	S/370™
Chipkill™	OS/390®	S/390®
DB2 Universal Database™	OS/400®	Seascape®
DB2®	PAL®	SecureWay®
DFS™	PartnerWorld®	StorWatch™
DirectTalk®	POWER Gt1™	System/370™
Domino™	POWER Gt3i™	System/390®
e-business on demand™	POWER Gt3™	SystemXtra®
Electronic Service Agent™	POWER Gt4e™	Tivoli®
Enterprise Storage Server®	POWER Gt4™	TME 10™
ESCON®	POWER Gt4x™	TME®
FICON™	POWER GTO™	TotalStorage®
FlashCopy®	POWER2™	TURBOWAYS®
GDPS®	POWER3™	Versatile Storage Server™
GXT1000™	POWER4+™	VisualAge®
GXT150L™	POWER4™	VM/ESA®
HACMP/6000™	POWERparallel®	VSE/ESA™
IBM®	PowerPC 601®	Wave®
IntelliStation®	PowerPC Architecture™	WebSphere®
iSeries™	PowerPC Reference Platform®	xSeries®
LoadLeveler®	PowerPC®	z/OS®
Lotus®	POWERserver®	zSeries®
Magstar®	pSeries®	
Micro Channel®	PTX®	

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

Preface

This IBM Redbook is a comprehensive single-source guide covering the most recent IBM eServer pSeries® and RS/6000® product line. In some cases, withdrawn products are included for comparison. Major hardware and software offerings are introduced and their prominent functions discussed.

This redbook is suitable for professionals wishing to acquire a better understanding of pSeries products, including:

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

Inside this redbook, you will find:

- ▶ A historical look at pSeries hardware
- ▶ An overview of the latest pSeries 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

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

This redbook does not replace the latest pSeries 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.

Scott Vetter is a Certified Project Manager at the International Technical Support Organization, Austin Center. He has enjoyed 20 years of rich and varied experience working for IBM in a variety of challenging roles. His latest efforts are directed at providing world-class pSeries and UNIX Redbooks™, whitepapers, and workshop collateral.

Flavio Cappelletti is an IBM @server pSeries systems product engineer at the EMEA Midrange Field Support Organization in Rome, Italy, supporting IBM sales, IBM Business Partners, IBM technical support organizations, and the IBM Dublin Enterprise Servers manufacturing plant. He is also a resolution assistant resource for critical situations related to RS/6000 and pSeries systems for the pSeries Customer Satisfaction Project Office. He has 27 years of IBM seniority and has 14 years of experience in the RS/6000 and pSeries systems.

JaeHoon Chung is an IT Specialist with the pSeries post-sales technical support team in IBM Korea. He has three years of experience working on pSeries and RS/6000 products. He is an IBM Certified Specialist in pSeries systems and AIX. He provides customers with technical support on pSeries systems, AIX, HACMP, and system management.

Rosario Delgado is a pSeries I/T Specialist with the Latin America Techline Team. She has six years of experience with the pSeries platform, having worked three years as a RS/6000 Telesales Specialist and three years in Techline. Her current area of expertise is technical marketing for pSeries systems.

Kimberly Fidler is an IT/Specialist with the pSeries Techline Team in Southfield, Michigan. She has more than seven years of experience in the I/T industry, the last three of which she has devoted to pSeries systems. She holds a B.A. degree from Michigan State University. She is an IBM Certified Specialist in pSeries Solutions Sales.

Hansjörg Schneider is a pSeries Hardware Support specialist at the IBM ITS Central Region Front Office in Ehningen, Germany. He has nine years of experience in the pSeries, RS/6000, and AIX fields. He has worked at IBM Germany for 15 years. He supports Customers and Customer Engineers to solve critical situations related to pSeries and RS/6000 systems. His areas of expertise include all pSeries Hardware, RS/6000 Hardware, and SSA Storage Systems, and he is also a AIX certified specialist.

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

Bruno Blanchard - IBM France

Robert Bluethman - IBM US

Anita Govindjee - IBM US

Volker Haug - IBM Germany

Adnan Ikram - IBM Pakistan

Keigo Matsubara - IBM Austin

Bill Mihaltse - IBM US

Basil Moshous - IBM Germany

Simon Robertson - IBM UK

Steve Schaffer - IBM US

Chris Smith - IBM New Zealand

Marcy Titus - IBM US

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:
redbook@us.ibm.com
- ▶ Mail your comments to:
IBM® Corporation, International Technical Support Organization
Dept. JN9B Building 003 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493



Introduction to IBM @server pSeries systems

IBM @server pSeries systems are powerful UNIX servers with mainframe-inspired reliability, availability, and autonomic computing features that create self-managing, resilient, responsive, efficient, and secure computing solutions. pSeries systems can manage diverse tasks ranging from engineering design to mission-critical applications, such as ERP, CRM, and Web serving tasks all the way up to massively parallel clustered high performance computing and business intelligence solutions. The pSeries family also combines leading-edge IBM technologies, including POWER4™, POWER4+ processors, and autonomic computing computer systems. Through high-performance and the flexibility to choose between AIX and Linux operating environments, pSeries delivers reliable, cost-effective solutions for commercial and technical computing applications in the entry, mid-range, and high-end UNIX segments.

The capability to perform dynamic logical partitioning (LPAR) further enhances the value of these servers.

In this chapter, we introduce a brief overview of the products, the processor architecture, and autonomic computing.

1.1 Overview of pSeries systems

The IBM @server brand was introduced in October 2000, replacing the RS/6000 brand first launched in February 1990. Since October 2000, new servers with UNIX operating systems have been introduced by the name of IBM @server pSeries systems.

pSeries systems come in various models, from tower servers to midrange to rack-optimized and large-scale systems. Figure 1-1 provides a general overview of the product family.




Figure 1-1 pSeries product line

All microprocessors in new pSeries servers make use of copper chip wiring, which offers 40 percent better conductivity than aluminum, improving chip performance and reducing power consumption. IBM's 64-bit POWER4 and POWER4+ microchips are manufactured with silicon-on-insulator (SOI) technology, which protects the millions of tiny transistors on a chip with a *blanket* of insulation, reducing harmful electrical effects that consume energy and hinder performance. Other forms of insulation may be used in the future. In addition, built-in intelligence features of the pSeries servers provide self-correcting capabilities that can minimize outages and keep applications running.

The pSeries platform addresses the need for reliability by providing high-availability solutions to meet today's requirements for e-business. To meet these requirements, the pSeries products offer a full range of high-performance

servers with a full set of highly functional state-of-the-art software to match the highest customer requirements of reliability, scalability, manageability, and security.

1.2 Processor descriptions

Though the entire system architecture contributes to the performance of the pSeries product line, the processors are a key component of system performance. The following sections outline the architectures of the latest pSeries microprocessors. Figure 1-2 groups the processors, showing their general evolution.




Figure 1-2 pSeries microprocessor history

1.2.1 32-bit versus 64-bit computing

64-bit computing is the current trend for all pSeries products. 64-bit microprocessors are faster than their 32-bit counterparts. They benefit from the latest design techniques and manufacturing processes, but there is more to it than that. 64-bit designs have inherent architectural advantages. These include 64-bit data flow, 64-bit arithmetic, and 64-bit addressing. Some workloads benefit from these advantages more than others. 64-bit data flow in modern RISC microprocessors are load/store machines. This means that all processing is done on data residing in registers. 64-bit microprocessors have 64-bit registers. 32-bit microprocessors have 32-bit registers. Data must flow between the registers and

memory to accomplish any operation. 64-bit microprocessors move 64 bits of data in the same amount of time it takes 32-bit microprocessors to move 32 bits of data. They can move data twice as fast as 32-bit microprocessors, but the data must be longer than 32 bits to take advantage of it.

The benefits of 64-bit architecture can be summarized as follows:

- ▶ Extended-Precision integer arithmetic
- ▶ Access to larger executables
- ▶ Access to larger data
- ▶ Access to larger file datasets
- ▶ Access to larger physical memory
- ▶ Access to higher SMP server scalability

1.2.2 POWER1

The first RS/6000 products were announced by IBM in February of 1990, and were based on a multiple chip implementation of the POWER architecture. This technology is now commonly referred to as POWER1, in the light of more recent developments. The models introduced included an 8 KB instruction cache (I-cache) and either a 32 KB or 64 KB data cache (D-cache). They had a single floating-point unit capable of issuing one compound floating-point multiply/add (FMA) operation each cycle, with a latency of only two cycles. Therefore, the peak MFLOPS rate was equal to twice the MHz rate. For example, the Model 530 was a desk-side workstation operating at 25 MHz, with a peak performance of 50 MFLOPS. Commonly occurring numerical kernels were able to achieve performance levels very close to this theoretical peak.

In January of 1992, the 7011 Model 220 was announced, based on a single chip implementation of the POWER architecture, usually referred to as RISC Single Chip (RSC). It was designed as a low-cost, entry-level desktop workstation, and contained a single 8 KB combined instruction and data cache.

The last POWER1 machine, announced in September of 1993, was the Model 580. It ran at 62.5 MHz and had a 32 KB I-cache and a 64 KB D-cache.

1.2.3 POWER2

Announced in September 1993, the first POWER2™ machines included the 55 MHz 7013 Model 580H, the 66.5 MHz Model 590, and the 71.5 MHz 990. The most significant improvement introduced with the POWER2 architecture for scientific and technical applications was the floating-point unit (FPU) that was enhanced to contain two 64-bit execution units. Thus, two floating-point

multiply/add instructions could be executed each cycle. A second fixed-point execution unit was also provided. In addition, several new hardware instructions were introduced with POWER2:

- ▶ Quad-word storage instructions. The quad-word load instruction moves two adjacent double-precision values into two adjacent floating-point registers.
- ▶ Hardware square root instruction.
- ▶ Floating-point to integer conversion instructions.

Although the Model 590 ran with only a marginally faster clock than the POWER1-based Model 580, the architectural improvements listed above, combined with a larger 256 KB D-cache size, enabled it to achieve far greater levels of performance.

In October 1996, IBM announced the RS/6000 7013 Model 595. This was the first machine to be based on the P2SC (POWER2 Super Chip) processor. As its name suggests, this was a single chip implementation of the POWER2 architecture, enabling the clock speed to be increased further. The Model 595 ran at 135 MHz, and the fastest P2SC processors, found in the Model 397 workstation and RS/6000 SP Thin4 nodes, ran at 160 MHz, with a theoretical peak speed of 640 MFLOPS.

1.2.4 PowerPC

The RS/6000 7011 Model 250 workstation, the first to be based on the PowerPC® 601® processor running at 66 MHz, was introduced in September 1993. The 601 was the first processor arising out of the partnership between IBM, Motorola, and Apple. The PowerPC architecture includes most of the POWER instructions. However, some instructions that were executed infrequently in practice were excluded from the architecture, and some new instructions and features were added, such as support for symmetric multiprocessor (SMP) systems. In fact, the 601 did not implement the full PowerPC instruction set, and was a bridge from POWER to the full PowerPC Architecture™ implemented in more recent processors, such as the 603, 604, and 604e. Currently, the fastest PowerPC-based machines from IBM for technical purposes, the four-way SMP system RS/6000 7025 Model F50 and the uniprocessor system RS/6000 43P 7043 Model 150, use the 604e processor running at 332 MHz and 375 MHz, respectively. The POWER3™ and POWER4 processors are also based on the PowerPC Architecture. IBM continues to develop the PowerPC chip and new applications are available in IBM and non-IBM products outside the pSeries family.

1.2.5 RS64

The first RS64 processor was introduced in September 1997 and was the first step into 64-bit computing for RS/6000. While the POWER2 product had strong floating-point performance, this series of products emphasized strong commercial server performance. It ran at 125 MHz with a 2-way associative, 4 MB L2 cache and had a 64 KB L1 instruction cache, a 64 KB L1 data cache, one floating-point unit, one load-store unit, and one integer unit. Systems were designed to use up to 12 processors. pSeries products using the RS64 were the first pSeries products to have the same processor and memory system as iSeries™ products.

In September 1998, the RS64-II was introduced. It was a different design from the RS64 and increased the clock frequency to 262 MHz. The L2 cache became 4-way set associative with an increase in size to 8 MB. It had a 64 KB L1 instruction cache, a 64 KB L1 data cache, one floating-point unit, one load-store unit, two integer units, and a short in-order pipeline optimized for conditional branches.

With the introduction of the RS64-III in the fall of 1999, this design was modified to use copper technology, achieving a clock frequency of 450 MHz, with a L1 instruction and data cache increased to 128 KB each. This product also introduced hardware multithreading for use by AIX. Systems were designed to use up to 24 processors.

In the fall of 2000, this design was enhanced to use silicon on insulator (SOI) technology, enabling the clock frequency to be increased to 600 MHz. The L2 cache size was increased to 16 MB on some models. Continued development of this design provided processors running at 750 MHz. The most recent version of this microprocessor was called the RS64-IV.

During the history of this family of products, industry leading performance were obtained for a large variety of benchmarks, including TPC-C (online transaction processing), SAP (enterprise resource planning (ERP)), Baan (ERP), PeopleSoft (ERP), SPECweb (Web serving), and SPECjbb (Java).

1.2.6 POWER3

The POWER3 processor brought together the fundamental design of the POWER2 micro architecture, as implemented in the P2SC processor, with the PowerPC architecture. It combined the excellent floating-point performance delivered by P2SC's two floating-point execution units, while being a 64-bit, SMP-enabled processor ultimately capable of running at much higher clock speeds than P2SC processors. Initially introduced in the fall of 1998 at a

processor clock frequency of 200 MHz, most recent versions of this microprocessor incorporate copper technology and operate at 450 MHz.

1.2.7 POWER4

The POWER4 system is a new generation of high-performance 64-bit microprocessors and associated subsystems especially designed for server and supercomputing applications. The following sections outline the architectures of the latest POWER4 microprocessor.

POWER4 system overview

The POWER4 processor is a high-performance microprocessor and storage subsystem utilizing IBM's most advanced semiconductor and packaging technology. A POWER4 system logically consists of multiple POWER4 microprocessors and a POWER4 storage subsystem, interconnected together to form an SMP system. Physically, there are three key components: the POWER4 processor chip, the L3 Merged Logic DRAM (MLD) chip, and the memory controller chip.

- ▶ The POWER4 processor chip contains two 64-bit microprocessors, a microprocessor interface controller unit, a 1.41 MB (1440 KB) level-2 (L2) cache, a level-3 (L3) cache directory, a fabric controller responsible for controlling the flow of data and controls on and off the chip, and chip/system pervasive functions.
- ▶ The L3 merged logic DRAM (MLD) chip, which contains 32 MB of L3 cache. An eight-way POWER4 SMP module will share 128 MB of L3 cache consisting of four modules, each of which contains two 16 MB merged logic DRAM chips.
- ▶ The memory controller chip features one or two memory data ports, each 16 bytes wide, and connects to the L3 MLD chip on one side and to the Synchronous Memory Interface (SMI) chips on the other.

POWER4 chip

The POWER4 chip is a result of advanced research technologies developed by IBM. Numerous technologies are incorporated into the POWER4 to create a high-performance, high-scalability chip design to power pSeries systems. Some of the advanced techniques used in the design and manufacturing processes of the POWER4 include copper interconnects and Silicon-on-Insulator.

Copper interconnects

As chips become smaller and faster, aluminum interconnects, which have been used in chip manufacturing for over 30 years, present increasing difficulties. In 1997, after nearly 15 years of research, IBM scientists announced a new

advance in the semiconductor process that involves replacing aluminum with copper. Copper has less resistance than aluminum, which permits the use of smaller circuits with reduced latency that allows for faster propagation of electrical signals. The reduced resistance and heat output make it possible to shrink the electronic devices even further while increasing clock speed and performance without resorting to exotic chip cooling methods.

Silicon-on-Insulator

Silicon-on-Insulator (SOI) refers to the process of implanting oxygen into a silicon wafer to create an insulating layer and using an annealing process until a thin layer of SOI film is formed. The transistors are then built on top of this thin layer of SOI. The SOI layer reduces the capacitance effects that consume energy, generate heat, and hinder performance.

Components

The components of the POWER4 chip are shown in Figure 1-3 on page 9. The chip has two processors on board. Included in what are referred to as the processor are the various execution units and the split first-level instruction and data caches. The two processors share a unified second level cache, also onboard the chip, through a Core Interface Unit (CIU). The CIU is a crossbar switch between the L2, implemented as three separate, autonomous cache controllers, and the two processors. Each L2 cache controller can operate concurrently and feed 32 bytes of data per cycle. The CIU connects each of the three L2 controllers to either the data cache or the instruction cache in either of the two processors. Additionally, the CIU accepts stores from the processors across 8-byte wide buses and sequences them to the L2 controllers. Each processor has associated with it a noncacheable (NC) Unit. The NC Unit in Figure 1-3 on page 9 is responsible for handling instruction serializing functions and performing any noncacheable operations in the storage hierarchy. Logically, this is part of the L2.




Figure 1-3 POWER4 chip logical view

The directory for a third level cache, L3, and (logically) its controller are also located on the POWER4 chip. The actual L3 is on a separate chip. A separate functional unit, referred to as the Fabric Controller, is responsible for controlling data flow between the L2 and L3 controller for the chip and for POWER4 communication. The GX controller is responsible for controlling the flow of information in and out of the system. Typically, this would be the interface to an I/O drawer attached to the system. But, with the POWER4 architecture, this is also where you can natively attach an interface to a switch for clustering multiple POWER4 nodes together.

Also included on the chip are functions we logically call Pervasive functions. These include trace and debug facilities used for First Failure Data Capture, Built-in Self Test (BIST) facilities, Performance Monitoring Unit, an interface to the Service Processor (SP) used to control the overall system, Power On Reset (POR) Sequencing logic, and Error Detection and Logging circuitry.

Four POWER4 chips can be packaged on a single module to form an 8-way SMP. Four such modules can be interconnected to form a 32-way SMP. To accomplish this, each chip has five primary interfaces. To communicate to other POWER4 chips on the same module, there are logically four 16-byte buses. Physically, these four buses are implemented with six buses, three on and three off. To communicate to POWER4 chips on other modules, there are two 8-byte buses, one on and one off. Each chip has its own interface to the off chip L3 across two 16-byte wide buses, one on and one off, operating at one third processor frequency. To communicate with I/O devices and other compute nodes, two 4-byte wide GX buses, one on and one off, operating at one third

processor frequency, are used. Finally, each chip has its own JTAG interface to the system service processor.

All of the buses previously discussed, except for the JTAG interface, scale with processor frequency. Over time, technological advances will allow an additional increase in processor frequency. As this occurs, bus frequencies will scale proportionately, allowing system balance to be maintained.

POWER4+ chip

POWER4+ is IBM's newest 64-bit microprocessor, which takes advantage of the most advanced 0.13 micron fabrication process and contains over 180 million transistors. The POWER4+ chip is available at speeds of 1.2, 1.45, 1.5, and 1.7 GHz. POWER4+ is based on POWER4 and also contains two processors, a high-bandwidth system switch, a large memory cache, and I/O interface.

L1 and L2 caches and L2 and L3 directories on the POWER4+ chip are manufactured with spare bits in their arrays that can be accessed using programmable steering logic to replace faulty bits in the respective arrays. This is analogous to the redundant bit steering employed in main store as a mechanism to avoid physical repair that is also implemented in POWER4+ systems. The steering logic is activated during processor initialization and is initiated by the Built-in system-test (BIST) at power on time. L3 cache redundancy is implemented at the cache line granularity level. Exceeding correctable error thresholds while running causes invocation of a dynamic L3 cache line delete function, capable of up to two deletes per cache. In the rare event of solid bit errors exceeding this quantity, the cache continues to run, but a message calling for deferred repair is issued. If the system is rebooted without such repair, the L3 cache is placed in bypass mode and the system comes up with this cache deconfigured.

1.3 Autonomic computing

The IBM autonomic computing initiative is about using technology to manage technology. This initiative is an ongoing effort to create servers that respond to unexpected capacity demands and system errors without human intervention. The goal is new highs in reliability, availability, and serviceability, and new lows in downtime and cost of ownership.

Today's pSeries offers some of the most advanced self-management features for UNIX servers on the market today. Autonomic computing on IBM @server pSeries servers describes the many self-configuring, self-healing, self-optimizing, and self-protecting features that are available on IBM @server pSeries servers.

Self-configuring

Autonomic computing provides self-configuration capabilities for the IT infrastructure. Today, IBM systems are designed to provide this at a feature level with capabilities such as plug-and-play devices and configuration setup wizards. Examples include:

- ▶ Virtual IP address (VIPA)
- ▶ IP multipath routing
- ▶ Microcode discovery services/inventory scout
- ▶ Hot-swappable disks
- ▶ Hot-plug PCI
- ▶ Wireless/pervasive system configuration
- ▶ TCP explicit congestion notification

Self-healing

For a system to be self-healing, it must be able to recover from a failing component by first detecting and isolating the failed component, taking it offline, fixing or isolating it, and reintroducing the fixed or replacement component into service without any application disruption. Examples include:

- ▶ Multiple default gateways
- ▶ Automatic system hang recovery
- ▶ Automatic dump analysis and e-mail forwarding
- ▶ EtherChannel automatic failover
- ▶ Graceful processor failure detection and failover
- ▶ First failure data capture
- ▶ Chipkill™ ECC Memory, dynamic bit-steering
- ▶ Memory scrubbing
- ▶ Automatic, dynamic deallocation (processors, PCI buses/slots)
- ▶ Electronic Service Agent™ (Call Home support)

Self-optimization

Self-optimization requires a system to efficiently maximize resource utilization to meet the end-user needs with no human intervention required. Examples include:

- ▶ Workload manager enhancement
- ▶ Extended memory allocator

- ▶ Reliable, scalable cluster technology (RSCT)
- ▶ PSSP cluster management and Cluster Systems Management (CSM)

Self-protecting

Self-protecting systems provide the ability to define and manage the access from users to all of the resources within the enterprise, protect against unauthorized resource access, detect intrusions and report these activities as they occur, and provide backup/recovery capabilities that are as secure as the original resource management systems. Examples include:

- ▶ Kerberos Version 5 authentication (authenticates requests for service in a network)
- ▶ Self-protecting kernel
- ▶ LDAP directory integration (LDAP aids in the location of network resources)
- ▶ SSL (manages Internet transmission security)
- ▶ Digital Certificates
- ▶ Encryption (prevents unauthorized use of data)

1.3.1 e-business on demand

In October 2002, IBM announced their vision of the next major phase of e-business adoption and called it e-business on demand™. In fact, e-business on demand is not just a vision, nor is it new. It is a statement of IBM's belief of how businesses will need to transform themselves to be successful. Businesses will have to adapt to cope with ever-increasing pressures from competition and other factors associated with the global economy. This implies a transformation to a fully integrated business across people, processes, and information, including suppliers and distributors, customers and employees.

IBM defines an on demand business as an enterprise whose business processes, integrated end-to-end across the company and with key partners, suppliers, and customers, can respond with speed to any customer demand, market opportunity, or external threat.

There are four key attributes of an on demand business:

- ▶ **Responsive:** Able to sense and respond to dynamic, unpredictable changes in demand, supply, pricing, labor, competition, capital markets, and the needs of its customers, partners, suppliers, and employees.
- ▶ **Variable:** Able to adapt processes and cost structures to reduce risk while maintaining high productivity and financial predictability.

- ▶ Focused: Able to concentrate on its core competencies and differentiating capabilities.
- ▶ Resilient: Able to manage changes and external threats while consistently meeting the needs of all of its constituents.

These attributes define the business itself. For a business to successfully attain and maintain these attributes, it must build an IT infrastructure that is designed to specifically support the business' goals. We call this infrastructure the on demand operating environment.

On demand operating environment

So what is an on demand operating environment? It is not a specific set of hardware and software. Rather, it is an environment that supports the needs of the business, allowing it to become and remain responsive, variable, focused, and resilient.

An on demand operating environment unlocks the value within the IT infrastructure to be applied to solving business problems. It is an integrated platform, based on open standards, to enable rapid deployment and integration of business applications and processes. Combined with an environment that allows true virtualization and automation of the infrastructure, it enables delivery of IT capability on demand.

An on demand operating environment must be:

- ▶ Flexible
- ▶ Self-managing
- ▶ Scalable
- ▶ Economical
- ▶ Resilient
- ▶ Based on open standards

IBM provides offerings that can be categorized into three primary areas, as shown in Figure 1-4 on page 14:

- ▶ Integration: Provides the facilities to gain a unified view of processes, people, information, and systems.
- ▶ Automation: Overcomes the complexity of systems management to enable better use of assets, improved availability and resiliency, and reduced operating costs.
- ▶ Virtualization: Simplifies deployment and improves use of computing resources by hiding the details of the underlying hardware and system

software, allowing for consolidation and the ability to adapt to changing demand.

The value of the operating environment is in the ability to dynamically link business processes and policies with the allocation of IT resources using offerings across all of these categories. In the operating environment, resources are allocated and managed without intervention, enabling resources to be used efficiently based on business requirements. Having flexible, dynamic business processes increases the ability to grow and manage change within the business.

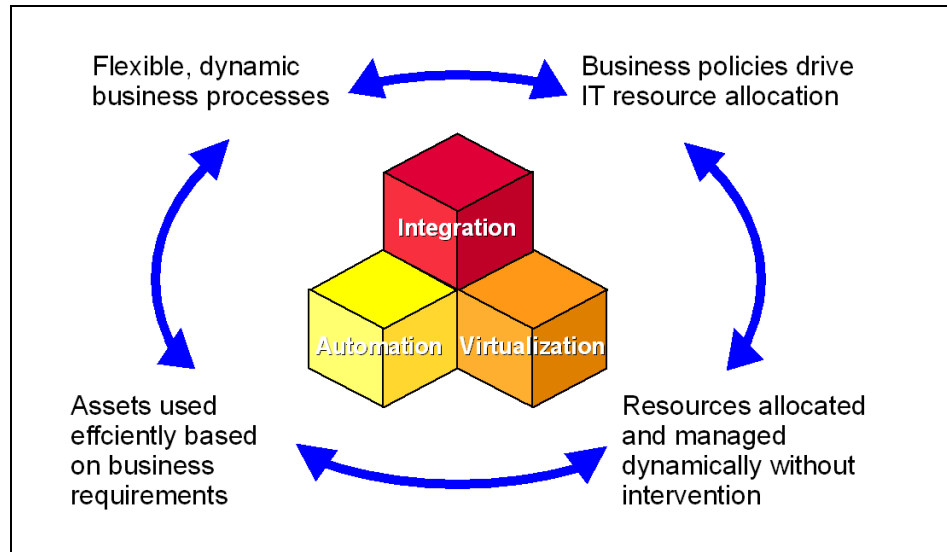


Figure 1-4 Overview of an on demand operating environment

1.3.2 Reliability, Availability, and Serviceability (RAS) features

Excellent quality and reliability are inherent in all aspects of the IBM @server pSeries design and manufacture, and the fundamental principle of the design approach is to minimize outages. The RAS features help to ensure that the system operates when required, performs reliably, and efficiently handles any failures that may occur. This is achieved using capabilities provided by both the hardware and the AIX 5L™ operating system.

Mainframe-class diagnostic capability based on internal error checkers, First-Failure Data Capture (FFDC), and run-time analysis are provided. This monitoring of all internal error check states is provided for processor, memory, I/O, power, and cooling components, and is aimed at eliminating the need to try to recreate failures later for diagnostic purposes. The unique IBM RAS capabilities are important for the availability of your server.

The following features provide the IBM @server pSeries with UNIX industry-leading RAS features:

- ▶ Fault avoidance through highly reliable component selection, component minimization, and error handling technology designed into the chips
- ▶ Improved reliability through processor operation at a lower voltage, enabled by the use of copper chip circuitry and Silicon-on-Insulator technology
- ▶ Fault tolerance through an additional hot-swappable power supply, and the capability to perform concurrent maintenance for power and cooling
- ▶ Automatic First Failure Data Capture (FFDC) and diagnostic fault isolation capabilities
- ▶ Concurrent run-time diagnostics based on First Failure Data Capture
- ▶ Predictive failure analysis on processors, cache, memory, and disk drives
- ▶ Dynamic error recovery
- ▶ Error Checking and Correction (ECC) or equivalent protection (such as refetch) on main storage, all cache levels (1, 2, and 3), and internal processor arrays
- ▶ Dynamic processor deallocation based on run-time errors (requiring more than one processor)
- ▶ Persistent processor deallocation (boot-time deallocation based on run-time errors)
- ▶ Persistent deallocation extended to memory
- ▶ Chipkill correction in memory
- ▶ Memory scrubbing and redundant bit-steering for self-healing
- ▶ Industry-leading PCI-X bus parity error recovery, as first introduced on the p690 systems
- ▶ Hot-plug functionality of the PCI-X bus I/O subsystem
- ▶ PCI-X bus and slot deallocation
- ▶ Disk drive fault tracking that monitors the number/rate of data errors and thresholds several recoverable hardware errors
- ▶ Avoiding checkstops with process error containment
- ▶ Environmental monitoring (temperature and power supply)
- ▶ Auto-reboot

1.3.3 Capacity Upgrade on Demand

Capacity Upgrade on Demand (CUoD) is available for pSeries 650, 670, and 690 with AIX 5L Version 5.2 and dynamic LPAR offers the capability to non-disruptively activate (no boot required) processors and memory. There is also the ability to temporarily activate processors to match intermittent performance needs. Combined with pSeries advanced technology, CUoD offers significant value for installations wanting to economically add new workloads on the same server or respond to increased workloads.

Pay as you grow

The CUoD option from IBM allows companies to install (spare or extra) processors and memory at an extremely attractive price and then bring new capacity online quickly and easily. With AIX 5L Version 5.2, processors and memory can be activated dynamically without interrupting system or partition operations.

CUoD processor options for pSeries 670 and 690 servers are available in units of four active and four inactive processors with up to 50 percent of the system in standby. p650 CUoD processors are available in pairs with a maximum of six in standby.

As workload demands require more processing power, unused processors can be activated in pairs simply by placing an order to activate the additional processors, sending current system configuration data to IBM, and receiving, over the Internet, an electronically encrypted activation key that unlocks the desired amount of processors. There is no hardware to ship and install, and no additional contract is required.

Memory activation works the same way. CUoD memory is available in various sizes for the p650, p670 and p690 systems. Activation in 4 GB increments is made by ordering an activation key to unlock the desired amount of memory.

On/Off Capacity on Demand (CoD)

For temporary workloads, pSeries offers an innovative solution with flexible processor activation. By ordering an On/Off CoD feature, the user receives an activation key that includes 60 days of temporary processor activation. Processors pairs can be then be turned on and off whenever needed. Charges are made against the 60 day processor allocation only when processors are activated. Increments of usage are measured in processor days and the minimum usage is one day per activated processor.

Trial Capacity on Demand

Trial Capacity on Demand enables CUoD features to be activated one time for a period of 30 consecutive days. If your system was ordered with CUoD features and they have not yet been activated, you can turn the features on for a one-time trial period. With the trial capability, you can gauge how much capacity you might need in the future, if you decide to permanently activate the resources you need.

Alternatively, the Trial Capacity on Demand function can be used to immediately activate resources while processing an order for a permanent activation code.

Capacity BackUp

Capacity Backup (CBU) is an on demand backup technology for high-end 16-way p670 and 32-way p690 servers. The servers, with On/Off Capacity on Demand capabilities, are similar to the iSeries for High Availability system introduced in September. IBM also has offered similar backup capabilities for its zSeries® mainframes.

The replicated p670 backup comes with twelve inactive and four active 1.45 GHz POWER4+ processors that can be activated if the production system goes down. The p690 is shipped with four POWER4+ processors active and another 28 inactive. Those chips can range in frequency from 1.3 GHz to 1.7 GHz.

Capacity BackUp systems are priced lower. If enterprises need to turn on inactive processors, they pay only for the power they use.

CUoD for the pSeries summary

For those customers who want to reduce their total cost of ownership, provide for fast, non-disruptive upgrades, and improve system availability and utilization, CUoD for the p650, p670, and the p690 is the way to go.

Benefits of CUoD include:

- ▶ Simple, dynamic activation of additional processors and memory
- ▶ Temporary activation of processors with On/Off CoD
- ▶ Automatic dynamic processor sparing
- ▶ Increased processor granularity
- ▶ 30-day trial period
- ▶ No commitment for future purchases
- ▶ No restriction on resale of system

The information at the following URL briefly explains the CuOD process:

<http://www.ibm.com/servers/eserver/pseries/cuod/>

Facts and features reference

The following section is taken from the brochure *Web Servers: IBM @server pSeries, IBM RS/6000, and IBM IntelliStation Facts and Features*. Figure 2-1 shows the RS/6000, IntelliStation®, and pSeries Models that are the targets of this redbook. An up-to-date version of this material is available at:

<http://www.ibm.com/servers/eserver/pseries/hardware/factsfeatures.html>

http://www.ibm.com/servers/eserver/pseries/hardware/workstations/facts_features.html

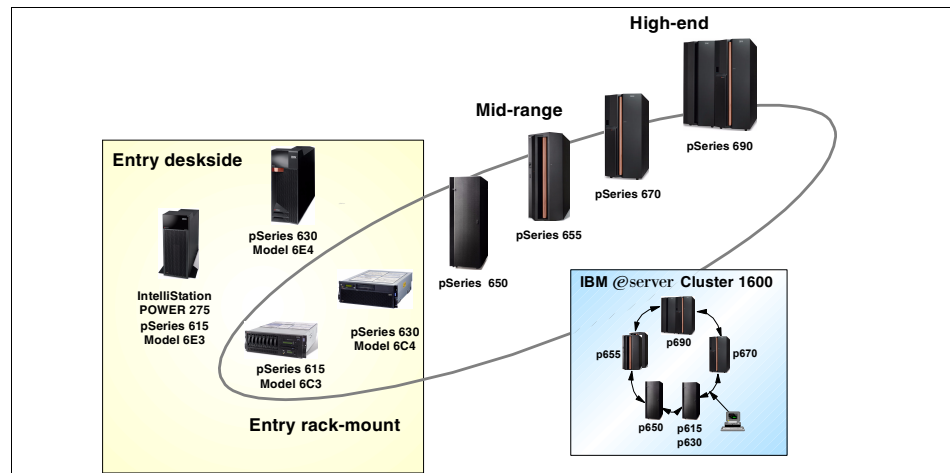


Figure 2-1 pSeries and IntelliStation products in focus

Table 2-1, Table 2-2 on page 21, Table 2-3 on page 22, and Table 2-4 on page 23 outline the important characteristics of the featured RS/6000, IntelliStation, and pSeries models.

Table 2-1 Facts and features for RS/6000 Model 170, and IntelliStation 265, 275

Product line	RS/6000	IntelliStation	IntelliStation POWER 275
Machine type-model	7044-170	9112-265	9114-275
System packaging	deskside	deskside	deskside
Microprocessor type	64-bit POWER3-II	64-bit POWER3-II	64-bit POWER4+
# of processors/system	1	1 or 2	1 or 2
Clock rates available	333/400/450 MHz	450 MHz	1.0 GHz ² /1.45 GHz
System memory (standard/maximum)	256 MB/2 GB	256 MB/8 GB ^a	1 GB-16 GB ^a
Data/instruction (L1) cache	64 KB/32 KB	64 KB/32 KB ^b	32 KB/64 KB ^b
Level 2 (L2) cache	1/4/8 MB	4 MB ^b	1.5 MB ^b
Level 3 (L3) cache	-	-	8 MB
Reliability, availability, serviceability			
Chipkill memory	-	-	X
Service Processor	X	X	X
Hot-swappable disks (internal and external)	-	X	X
Dynamic Processor Deallocation	-	-	X
Hot-plug slots	-	-	-
Redundant hot-plug power	-	O	O
Redundant hot-plug cooling	-	X	X
Capacity			
On demand features	-	-	-
Logical partitions (maximum)	-	-	-
Slots available (maximum)	6 PCI (4 32-bit + 2 64-bit)	5 PCI (2 32-bit+ 3 64-bit)	6 (2 32-bit, 4 64-bit)
PCI bus speed (maximum)	50 MHz	50 MHz	133 MHz
Disk/media bays	3/3	7/3	4/3
Minimum/maximum internal disk	18.2 GB/293.6 GB	18.2 GB/1.0 TB	36.4 GB/587.2 GB
Ethernet ports	1 10/100	1 10/100	1 - 10/100;1 - 10/100/1000
Storage interfaces			
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra3 SCSI	X	X	X
PCI 4-Channel Ultra3 SCSI RAID	X	X	X
PCI-X Dual Channel Ultra320 SCSI	-	-	X
SSA Advanced SerialRAID Plus	X	X	X
2 Gigabit Fibre Channel PCI	X	X	-
2 Gigabit Fibre Channel PCI-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
10/100 Mbps Ethernet PCI II	X	X	X
10/100/1000 Mbps Base-TX Ethernet PCI-X	X	X	X
Gigabit Ethernet (Fibre I UTP)	X/X	X/X	X/X
Gigabit Ethernet-SX PCI-X	X	X	X
ATM 155 Mbps (Fibre I UTP)	X/X	X/X	X/X
ATM 622 Mbps (Fibre)	X	X	X
HIPP ^k	-	-	X
Digital Trunk Quad ^k	X	-	X
X.25 ^k	-	-	X
SDLC	X	-	X
BSC	X	-	-
e-business Cryptographic Accelerator.	X	X	X
PCI Cryptographic Coprocessor (FIPS-4)	X	-	-

Product line	RS/6000	IntelliStation	IntelliStation POWER 275
Graphics accelerators	GXT135P	GXT135P	GXT135P, GXT4500P, GXT6500P

Table 2-2 Facts and features for pSeries Model 6C1, 6E1, 6C3, 6E3, and 6E4

Product line	pSeries 610	pSeries 615	pSeries 630
Machine type-model	7028-6C1/6E1	7029-6C3/6E3	7028-6E4
System packaging	19" rack drawer-5U/ deskside	19" rack drawer-4U/ deskside	deskside
Microprocessor type	64-bit POWER3-II	64-bit POWER4+	64-bit POWER4+
# of processors/system	1 or 2	1 or 2	1, 2, or 4
Clock rates available (standard/option)	333/375/450 MHz	1.2 GHz/1.45 GHz ^l	1.2 GHz/1.45 GHz
System memory (standard/maximum)	512 MB/8 GB ^a	1 GB/16 GB ^a	1 GB/32 GB ^{a,1}
Data/instruction (L1) cache	64 KB/32 KB ^b	32KB/64 KB ^b	32 KB/64 KB ^b
Level 2 (L2) cache	4/4/8 MB ^b	1.5 MB ^g	1.5 MB ^g
Level 3 (L3) cache	-	8 MB ^g	8 MB ^g
Reliability, availability, serviceability			
Chipkill memory	-	X	X
Service Processor	X	X	X
Hot-swappable disks (internal and external)	X ^s	X	X
Dynamic Processor Deallocation	X	X	X
Dynamic deallocation: PCI bus slot	-	X	X
Hot-plug slots	-	X	X
Redundant hot-plug power	O	O	O
Redundant hot-plug cooling	X	X	X ^f
NEBS3	-	-	-
Capacity			
On demand features	-	-	-
Logical partitions (maximum)	-	-	3
Slots available (maximum)	5 PCI (two 32-bit + three 64-bit)	6 PCI (two 32-bit + four 64-bit)	6 PCI (64-bit)
PCI bus speed (maximum)	50 MHz	133 MHz	133 MHz
Disk/media bays	7 ^{s/3}	8/3	4/2
Minimum/maximum internal disk	36.4 GB/1.0 TB	36.4 GB/1.1 TB	36.4 GB/587.2 GB
Required/optional I/O drawers	0/0	0/0	-
Storage interfaces			
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra3 SCSI	X	-	-
PCI-X 2-Channel Ultra320 SCSI	X	X	X
PCI-X 2-Channel Ultra320 RAID	X	X	X
SSA Advanced SerialRAID Plus	X	X	X
2 Gigabit Fibre Channel PCI	X	-	X
2 Gigabit Fibre Channel PCI-X	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 PCI II	X	X	X
10/100/1000 mbps Base-TX Ethernet PCI-X	X	X	X
Gigabit Ethernet Fibre I UTP)	X/X	X/X	X/X
Gigabit Ethernet-SX PCI-X	X	X	X
ATM 155 Mbps (Fibre I UTP)	X/X	X/X	X/X
ATM 622 Mbps (Fibre)	X	X	X
SP System Attachment	-	-	-
ESCON® Control Unit (host attach) ^k	X ^c	-	-
ESCON Emulation (tape attach) ^k	X ^c	-	-
HIPPI ^k	X ^c	-	X
Digital Trunk Quad ^k	X	X	X

Product line	pSeries 610	pSeries 615	pSeries 630
SDLC		X	X
BSC		X	X
e-business Cryptographic Accelerator.	X	X	X
PCI Cryptographic Coprocessor (FIPS-4)	X	X	X
Display adapter	GXT135P	GXT135P	GXT135P, GXT4500P, GXT6500P

Table 2-3 Facts and features for pSeries Models 6C4, B80, and 6M2

Product line	pSeries 630	pSeries 640	pSeries 650
Machine type-model	7026-6C4	7026-B80	7038-6M2
System packaging	19-in. rack drawer-4U	19-in. rack drawer-5U	19-in. rack drawer-8U system/4U - I/O ^{bb}
Microprocessor type	64-bit POWER4+	64-bit POWER3-II	64-bit POWER4+
# of processors/system	1, 2, or 4	1, 2, 3, or 4	2, 4, 6, or 8 ^{aa}
Clock rates available (standard/option)	1.2/1.45 GHz	375/450 MHz	1.2/1.45 GHz
System memory (standard/maximum)	1 GB/32 GB ^{a, j}	256 MB/16 GB ^a	2 GB/64 GB ^a
Data/instruction (L1) cache	32 KB/64 KB ^b	64 KB/32 KB ^b	32 KB/64 KB ^b
Level 2 (L2) cache	1.5 MB ^g	4/8 MB ^b	1.5 MB ^g
Level 3 (L3) cache	8 MB ^g	-	8/32 MB ^g
Reliability, availability, serviceability			
Chipkill memory	X	-	X
Service Processor	X	X	X
Hot-swappable disks (internal and external)	X	X	X
Dynamic Processor Deallocation	X	X	X
Dynamic deallocation: PCI bus slot	X	-	X
Hot-plug slots	X	-	X
Redundant hot-plug power	O	O	X
Redundant hot-plug cooling	X ^f	X	X
NEBS3	X	X	-
Capacity			
On demand features	-	-	P, M, O/O ^f
Logical partitions (maximum)	4 ^{dd}	-	8
Slots available (maximum)	20 PCI (64-bit) ^{dd}	5 PCI (one 32-bit + four 64-bit)	63 PCI (one 32-bit + 62 64-bit) ^{dd}
PCI bus speed (maximum)	133 MHz	50 MHz	133 MHz
Disk/media bays	28 ^{e/2}	5	100 ^{e/2}
Minimum/maximum internal disk	36.4 GB/4.1 GB ^{aa}	36.4 GB/293.6 GB	36.4 GB/14.6 TB ^e
Required/optional I/O drawers	0/2	0/0	0/8
Storage interfaces			
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra3 SCSI	-	X	-
PCI-X 2-Channel Ultra320 SCSI	X	X	X
PCI-X 2-Channel Ultra320 RAID	X	X	X
SSA Advanced SerialRAID Plus	X	X	X
2 Gigabit Fibre Channel	X	X	X
2 Gigabit Fibre Channel PCI-X	X	X	X
Communications and connectivity			
EIA RS232D	X	X	X
Token ring 4/16 Mbps	X	X	X
4-Port 10/100 Mbps Ethernet	X	X	X
Ethernet 10/100 Mbps PCI II	X	X	X
10/100/1000 mbps Base-TX Ethernet PCI-X	X	X	X
Gigabit Ethernet Fibre I UTP)	X/X	X/X	X/X
Gigabit Ethernet-SX PCI-X	X	X	X
ATM 155 Mbps (Fibre I UTP)	X/X	X/X	X/X
ATM 622 Mbps (Fibre)	X	X	X
2 Gigabit Fibre Channel PCI-X	X	X	X
SP System Attachment	X	-	X

Product line	pSeries 630	pSeries 640	pSeries 650
ESCON Control Unit/Emulation	-	-	X
HIPPJ ^k	X	-	X ^h
Digital Trunk Quad ^k	X	X	X
SDLC	X	X	X
BSC	X	X	X
e-business Cryptographic Accelerator	X	X	X
PCI Cryptographic Coprocessor	X	X	X
Display adapter	GXT135P	GXT135P, GXT4500P, GXT6500P	GXT135P

Table 2-4 Facts and features for pSeries Models 655, 760, and 6901

Product line	pSeries 655	pSeries 670	pSeries 690
Machine type-model	7039-651	7040-671	7040-681
System packaging	24-in. system frame	24-in. system frame	24-in. system frame (+ I/O frame)
Microprocessor Type	64-bit POWER4/4+	64-bit POWER4/4+	64-bit POWER4/4+
# of processors/system	4 or 8 ^{ff}	4 ^{gg} , 8, or 16 ^{aa}	8, 16, 24, or 32 ^{aa}
Clock rates available (standard/option)	1.1/1.3/1.5/1.7 GHz	1.1/1.5 GHz	1.1/1.3/1.5/1.7 GHz
System memory (standard/maximum)	4 GB/32 GB ^{a,m}	8 GB/256 GB ^{a,m}	8 GB/512 GB ^a
Data/instruction (L1) cache	32 KB/64 KB ^b	32 KB/64 KB ^b	32 KB/64 KB ^b
Level 2 (L2) cache	5.7/6.0 MB ^g	5.7/6.0 MB ^g	5.7/6.0 MB ^g
Level 3 (L3) cache	128 MB ^g	128 MB ^g	128 MB ^g
Reliability, availability, serviceability			
Chipkill memory	X	X	X
Service Processor	X	X	X
Hot-swappable disks (internal and external)	X	X	X
Dynamic Processor Deallocation	X	X	X
Dynamic deallocation: PCI bus slot	X	X	X
Hot-plug slots	X	X	X
Redundant hot-plug power	X ^{cc}	X	X
Redundant hot-plug cooling	X ^{cc}	X	X
NEBS3	-	-	-
Capacity			
On demand features	-	P, B, M, O/O	P, B, M, O/O ^r
Logical partitions (maximum)	4	16	32
Slots available (maximum)	23 PCI (64-bit) ^{dd}	60 PCI (64-bit) ^{t,dd}	160 PCI (64-bit) ^{dd}
PCI bus speed (maximum)	133 MHz	133 MHz	133 MHz
Disk/media bays	18 ^e /-	48 ^e /5	160 ^e /5
Minimum/maximum internal disk	36.4 GB/2.3 TB ^e	36.4 GB/7.0TB ^e	36.4 GB/18.7 TB ^e
Required/optional I/O drawers	0/1	1/2	1/7
Storage interfaces			
Ultra SCSI SE and Ultra SCSI Differential	X	X	X
PCI 2-Channel Ultra3 SCSI	-	-	-
PCI 2-Channel Ultra320 SCSI	X	X	X
PCI-X 2-Channel Ultra320 RAID	X	X	X
SSA Advanced SerialRAID Plus	X	X	X
2 Gigabit Fibre Channel	X	X	X
2 Gigabit Fibre Channel PCI-X	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 PCI II	X	X	X
10/100/1000 mbps Base-TX Ethernet PCI-X	X	X	X
Gigabit Ethernet Fibre I UTP)	X/X	X/X	X/X
Gigabit Ethernet-SX PCI-X	X	X	X
ATM 155 Mbps (Fibre I UTP)	X/X	X/X	X/X
ATM 622 Mbps (Fibre)	X	X	

Product line	pSeries 655	pSeries 670	pSeries 690
2 Gigabit Fibre Channel PCI-X	X	X	X
SP System Attachment	X	X	X
ESCON Control Unit/Emulation ^k	X	X	X
HIPPI ^k	X	X	X
Digital Trunk Quad ^k	-	-	-
SDLC	X	X	X
BSC	X	X	X
e-business Cryptographic Accelerator.	X	X	X
PCI Cryptographic Coprocessor (FIPS-4)	X	X	X
Display adapter	GXT135P	GXT135P	GXT135P

Notes:

- ▶ X = Standard; Supported
- ▶ O = Optionally Available; Supported
- ▶ N/A = Not Applicable
- ▶ P = Processor Capacity Upgrade on Demand offering available
- ▶ B = pSeries for Capacity Backup offering
- ▶ M = Memory Capacity Upgrade on Demand offering available
- ▶ O/O = On/Off Capacity on Demand offering available
- ▶ SOD = Statement of General Direction announced
- ▶ ^a Shared memory
- ▶ ^b Per processor
- ▶ ^c Available only as a MES
- ▶ ^d 64 GB available via RPQ
- ▶ ^e Includes CEC and required I/O drawer components
- ▶ ^f Optional on 1.2 GHz systems
- ▶ ^g Per processor card or multichip module
- ▶ ^h Third disk bay available via RPQ
- ▶ ⁱ Available only as a 1-way system
- ▶ ^j 16 GB maximum on 1-way system
- ▶ ^k Requires additional software; check on availability
- ▶ ^l Available on a 2-way system only
- ▶ ^m 4 GB/128 GB for 4-way and 8-way systems
- ▶ ⁿ 256 MB/16 GB on 1-way system
- ▶ ^o Using 375 MHz processors and 4 MB of L2 cache
- ▶ ^p Using 450 MHz processors and 8 MB of L2 cache
- ▶ ^r Using 1.45 GHz processors
- ▶ ^s Requires an optional feature on 333 MHz systems
- ▶ ^t Reduced by 1/3 on 4-way and 8-way systems only
- ▶ ^u Supported on 1-way and 2-way systems only
- ▶ ^v Using 1.2 GHz processors
- ▶ ^w Using 1.3 GHz processors
- ▶ ^x Using 1.3 GHz HPC processors
- ▶ ^y Using 1.5 GHz processors
- ▶ ^z Using 1.7 GHz processors
- ▶ ^{aa} Capacity Upgrade on Demand (CUoD) supported

- ▶ ^{bb} Two 7311-D10 I/O drawers or one 7311-D20 I/O drawer per 4U of rack space
- ▶ ^{cc} Redundancy at the system frame level
- ▶ ^{dd} Using I/O drawer(s)
- ▶ ^{ee} First rPerf value - using one 2-way processor card; second rPerf value – using two 1-way processor cards
- ▶ ^{ff} 1.3/1.7 GHz systems are 4-way; 1.1/1.5 GHz systems are 8-way
- ▶ ^{gg} Using 1.1 GHz processors

Table 2-5 and Table 2-6 on page 26 show pSeries and IntelliStation unit details.

Table 2-5 System unit details

System unit details	RS/6000 150	RS/6000 170	pSeries 610	pSeries 615	pSeries 640	pSeries 630
Standard internal disk bays	2	3	8	4	4	4
Optional internal disk bays	-	-	-	4	-	-
Available media bays	1	1	1	3	1	2
- Standard size (xx-in height)	1	1	1	1	1	2
- Slimline size (yy-in height)	-	-	-	2	-	-
Standard diskette drive	X	X	X	-	-	-
Standard CD-ROM or DVD-ROM	X	X	X	-	-	-
RS232 Serial ports	2	2	3	3	3	3
Keyboard/Mouse Port	X	X	X	X	X	X
Parallel port	X	X	X	X	X	X
Integrated 10/100 Ethernet port	1	1	2	1	2	2
Integr. 10/100/1000 Ethernet port	-	-	-	1	-	-
PCI slots	5	6	5	-	5	-
- Short 32-bit 5V 33 MHz	3	4	2	-	1	-
- Long 32-bit 5V 33 MHz	2	-	-	-	-	-
- Long 64-bit 5V 33 MHz	-	-	1	-	2	-
- Long 64-bit 3.3V 50/66 MHz	-	2	2	-	2	-
PCI-X slots	-	-	-	6	-	6
- Long 64-bit 3.3V 133 MHz	-	-	-	4	-	6
- Short 64-bit 3.3V 133 MHz	-	-	-	-	-	-
- Short 32-bit 3.3V 133 MHz	-	-	-	-	-	-
- Short 32-bit 3.3V 66 MHz	-	-	-	2	-	-
RJ-4x connector**	-	-	RJ-48	RJ-48	RJ-45	RJ-48
Rack Light Indicator**	-	-	X	X	X	X
LED diagnostics	X	X	X	X	X	X
HMC ports	-	-	-	2	-	2
RIO ports**	-	-	-	-	-	2
RIO-2 ports**	-	-	-	-	-	2

Table 2-6 System unit details

System unit details	pSeries 650	pSeries 655	pSeries 670	pSeries 690	IntelliSt. 265	IntelliSt. 275
Standard internal disk bays	4	2	16	16	7	4
Optional internal disk bays	-	-	32	144	-	4
Available media bays	1	0	3	3	1	3
- Standard size (xx inch height)	1	-	3	3	1	1
- Slimline size (yy inch height)	-	-	-	-	-	2
Standard diskette drive	X	-	X	X	X	X
Standard CD-ROM or DVD-ROM	X	-	X	X	X	X
RS232 Serial ports	4	-	2	2	3	3
Keyboard/Mouse Port	X	-	X	X	X	X
Parallel port	-	-	-	-	X	X
Integrated 10/100 Ethernet port	1	2	-	-	2	2
Integr. 10/100/1000 Ethernet port	-	-	-	-	-	-
PCI slots	-	-	20	20	5	-
- Short 32-bit 5V 33 MHz	-	-	-	-	3	-
- Long 32-bit 5V 33 MHz	-	-	-	-	-	-
- Long 64-bit 5V 33 MHz	-	-	6	6	1	-
- Long 64-bit 3.3V 50/66 MHz	-	-	14	14	2	-
PCI-X slots	7	3	20	20	-	6
- Long 64-bit 3.3V 133 MHz	6	3	20	20	-	4
- Short 64-bit 3.3V 133 MHz	-	-	-	-	-	-
- Short 32-bit 3.3V 133 MHz	1	-	-	-	-	-
- Short 32-bit 3.3V 66 MHz	-	-	-	-	-	2
RJ-4x connector**	-	-	-	-	RJ-48	RJ-48
Rack Light Indicator**	X	-	-	-	X	X
LED diagnostics	X	X	X	X	X	X
HMC ports	2	2	2	2	-	2
RIO ports**	8	1	3	8	-	-
RIO-2 ports**	8	1	3	8	-	-

Notes:

- ▶ X = Available
- ▶ Rack-mount model only:
 - * = Assuming single required I/O drawer; either PCI or PCI-X slot drawer may be installed
 - ** = Rack-mount model only

Table 2-7 on page 27 shows the server I/O attachments for pSeries 630 (Model 6C4), p650, p655, p670, and p690.

Table 2-7 Server I/O attachment

Server I/O attachment	Max. drawer per system	Slots per drawer	Max. slots per system	Disk bays per drawer	Max. disk bays per system	Max. I/O drawer disk capacity	Max. disk capacity per system
pSeries 630 Model 6C4			20		28		4.1 TB
7311-D20 drawer ⁴	2	7 PCI-X		12		3.5 TB	
pSeries 650 ²			63		100		14.6 TB
7311-D10 drawer	8	1 PCI, 5 PCI-X					
7311-D20 drawer ⁴	8	7 PCI-X		12		14.0 TB	
pSeries 655			23		18		2.6 TB
7040-61D drawer ⁵	1	20 PCI or 20 PCI-X		16		2.3 TB	
pSeries 670 ¹			60 ³		48		7.0 TB
7040-61D drawer ⁵	3 ¹	20 PCI or 20 PCI-X		16		7.0 TB	
pSeries 690 ¹			160		128		18.7 TB
7040-61D drawer ⁵	8 ¹	20 PCI or 20 PCI-X		16		18.7 TB	

Notes:

- ▶ ¹ At least one drawer is required
- ▶ ² A maximum of eight drawers of any type may be installed
- ▶ ³ Reduced by 1/3 on 4- and 8-way systems
- ▶ ⁴ Ultra320 disk drives enabled
- ▶ ⁵ Ultra320 SCSI adapter provides access to external Ultra320 disk drives

Table 2-8 shows the peak bandwidth for pSeries 630, 650, 655, 670, and 690.

Table 2-8 Peak bandwidth

Peak bandwidth	pSeries 615	pSeries 630	pSeries 650	pSeries 655	pSeries 670	pSeries 690
Memory to L3 cache (GBps)	6.4	12.8	25.6	51.2	102.4	204.8
I/O (GBps) - RIO-2 ports	2	4	16	4	14	44

Table 2-9 and Table 2-10 report the standard warranty for RS/6000, pSeries, and IntelliStation units.

Table 2-9 Standard warranty in United States (other countries may vary)

Standard warranty	RS/6000 150	RS/6000 170	pSeries 610	pSeries 615	pSeries 640	pSeries 630
24 x 7 with two response objectives (some cities)	-	-	-	-	-	-
24 x 7 with four response objectives	X	X	-	-	X	-
9 x 5 next business day	-	-	X	-	-	X
9 x 5 next business day with Customer Replaceable units (CRU)	-	-	-	X	-	-

Table 2-10 Standard warranty in United States (other countries may vary)

Standard warranty	pSeries 650	pSeries 655	pSeries 670	pSeries 690	IntelliSt. 265	IntelliSt. 275
24 x 7 with two response objectives	-	-	X	X	X	-
24 x 7 with four response objectives	-	-	-	-	-	-
9 x 5 next business day	X	X	-	-	-	-
9 x 5 next business day with Customer Replaceable units (CRU)	-	-	-	-	-	X

Table 2-11 and Table 2-12 report the software support for RS/6000, pSeries, and IntelliStation units

Table 2-11 System software: Model 150 through Model 640

Operating system support	RS/6000 150	RS/6000 170	pSeries 610	pSeries 615	pSeries 630	pSeries 640
AIX 5L Version 5.1 (5765-E61)	X	X	X	X	X	X
AIX 5L Version 5.1 LPAR	-	-	-	-	X	-
AIX 5L Version 5.2 (5765-E61)	X	X	X	X	X	X
AIX 5L Version 5.2 dynamic LPAR	-	-	-	-	X	-
Linux distribution (64-bit) ¹	-	X	X	X	X	X
Linux LPAR ¹	-	-	-	-	X	-
HACMP 5.1 (5765-F62) ²	X	X	X	-	X	X
CSM for AIX 5L Version 1.3.1 (5765-F67)	-	-	-	X	X	-
PSSP Version 3.5 (5765-D51) ²	-	-	-	-	X	-
CSM for Linux on pSeries Version 1.3.2	-	-	-	X	X	-

Table 2-12 System software: Model 650 through Model 265

Operating system support	pSeries 650	pSeries 655	pSeries 670	pSeries 690	IntelliSt. 265	IntelliSt. 275
AIX 5L Version 5.1 (5765-E61)	X	X	X	X	X	X
AIX 5L Version 5.1 LPAR	X	X	X	X	-	-
AIX 5L Version 5.2 (5765-E61)	X	X	X	X	X	X
AIX 5L Version 5.2 dynamic LPAR	X	X	X	X	-	-
Linux distribution (64-bit) ¹	X	X ³	-	-	-	-
Linux LPAR ¹	X	X ⁴	X	X	-	-
HACMP 5.1 (5765-F62) ²	X	X	X	X	-	-
CSM for AIX 5L Version 1.3.1 (5765-F67)	X	X	X	X	-	-
PSSP Version 3.5 (5765-D51) ²	X	X	X	X	-	-

Operating system support	pSeries 650	pSeries 655	pSeries 670	pSeries 690	IntelliSt. 265	IntelliSt. 275
CSM for Linux on pSeries 1.3.2	X	X	-	-	-	-

Notes:

- ▶ X = Available or standard feature.
- ▶ SOD = Statement of General Direction announced.
- ▶ ¹ Many of the features described in this document are operating system dependent and may not be available on Linux. For more information, check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html
- ▶ ² Support on AIX systems only.
- ▶ ³ Planned for 2H 2003 availability from one or more Linux distributors.
- ▶ ⁴ Supported only when AIX is running in a separate LPAR.



Entry systems

This chapter provides an introduction and overview to the following models:

- ▶ 7028 Model 6C1 and 6E1 IBM @server pSeries 610
- ▶ 7029 Model 6C3 and 6E3 IBM @server pSeries 615
- ▶ 7028 Model 6C4 and 6E4 IBM @server pSeries 630
- ▶ 9114 Model 275 IntelliStation
- ▶ 7043 Model 150 IBM RS/6000
- ▶ 7044 Model 170 IBM RS/6000
- ▶ 9112 Model 265 IntelliStation
- ▶ 7026 Model B80 IBM @server pSeries 640

3.1 7028 Model 6C1 and 6E1 IBM @server pSeries 610

The IBM @server pSeries 610 Models 6C1 and 6E1 (referred to hereafter as the Model 6C1 and Model 6E1) are members of the 64-bit family of symmetric multiprocessing (SMP) UNIX servers from IBM and use 64-bit, copper-based, POWER3-II microprocessors. The Model 6C1 (product number 7028-6C1) is a rack-mounted server, while the Model 6E1 (7028-6E1) is a tower server.

Figure 3-1 shows the pSeries 610 Model 6C1 and Model 6E1.



Figure 3-1 p610 Model 6C1 and 6E1

Both models provide outstanding performance for many commercial and technical computing requirements, especially e-business, application or service providers, and database servers. Powered by the POWER3-II 64-bit processors, the Model 6C1 and Model 6E1 will bring significant value to those customers desiring a server solution for their e-business requirements and incorporating the power of the Web into the operations of their organizations. The availability of internal SCSI RAID features extends the use of these servers into application areas where cost, footprint, and reliability are all important factors.

3.1.1 Product positioning

With its SMP capability, the Model 6E1 and Model 6C1 are well positioned for transaction-intensive server solutions such as e-commerce, database, and application database server applications. The Model 6E1 and Model 6C1 provide 2-way SMP capability, up to 8 GB of memory, and up to 1.0 TB internal DASD, which makes it an excellent solution for many commercial applications. It is a cost-effective performance alternative to the RS/6000 7025 Models F50 or F80,

or 7026 Models H70 or H80 for commercial customers who do not require the expansion capabilities that these models offer.

Figure 3-2 shows both pSeries 610 Model 6C1 rack, and 6E1 desktop versions.



Figure 3-2 p610 Model 6C1 and 6E1 general view

3.1.2 Highlights

The pSeries 610 Model 6C1 and 6E1 deliver a cost-efficient growth path to the future through such attributes as:

- ▶ Outstanding value
- ▶ Powerful, symmetric multiprocessing (SMP) e-business server
- ▶ Rack-mount (Model 6C1) configuration
- ▶ Desktop (Model 6E1) configuration
- ▶ Up to two state-of-the-art IBM 64-bit POWER3-II processors that exploit unique copper-based technology for high performance and reliable service with the following processor options:
 - 1-way 333 MHz processor, including 4 MB of Level 2 (L2) ECC cache.

- 1-way 375 MHz processor, including 4 MB of L2 ECC cache.
- 1-way 450 MHz processor, including 8 MB of L2 ECC cache.
- ▶ Performance and capacity needed by demanding e-business applications
- ▶ Up to 8 GB of memory
- ▶ Up to 1.0 TB of internal disk storage
- ▶ Ten bays and five PCI slots for feature expandability
- ▶ An integrated service processor
- ▶ Hot-swap redundant power supply (AC)
- ▶ Processor deallocation
- ▶ Hot-swap cooling fans and disk drives
- ▶ A programmable beacon that can be triggered remotely to quickly alert an on-site technician to the system that needs maintenance or repair (Model 6C1 only)
- ▶ Light Path diagnostics

3.1.3 Technical overview

This section provides the technical overview of the p610 Model 6C1 and 6E1.

The Model 6C1 is a 5U (EIA) 19-in. rack-mounted system and has a size of 426 mm W x 617 mm D x 215 mm H (16.8 inches W x 24.0 inches D x 8.5 inches). The Model 6E1 is a tower package. Its size without option accessories, is 215 mm W x 617 mm D x 426 mm H (8.5 inches W x 24.0 inches D x 16.8 inches H). Both systems have a maximum weight of 43.1 kg (94.8 pounds).

Figure 3-3 on page 35 shows the high level system block diagram of both models.




Figure 3-3 p610 Model 6C1 and 6E1 system block diagram

The pSeries 610 Models 6C1 and 6E1 have two processor card slots and can accommodate three different processor cards: a 1-way 333 MHz, a 1-way 375 MHz, or a 1-way 450 MHz. Note that slot 1 must have a processor card installed for normal operation.

These models use a 64 KB data and a 32 KB instruction 128-way set associative L1 cache. The size of both data and instruction cache reduces the number of cache misses, results in more cache hits, and maximizes performance. Both data and instruction cache are parity protected.

The L1 cache is effectively supplemented by a 4 MB 4-way set associative L2 cache, which is located on the 333 MHz and 375 MHz processor cards (8 MB for the 450 MHz processor card). The speed of the L2 cache is dependent upon the processor speed. The POWER3-II uses a private 32-byte L2 cache bus, operated at 166.5 MHz with the 333 MHz processor (1:2 ratio) 250 MHz with the 375 MHz processor card (2:3 ratio), and 225 MHz with the 450 MHz processor card (1:2 ratio). Both the enhanced clock speed and 4-way set associative L2 cache improve cache efficiency. The L2 controller uses a least recently used (LRU) algorithm to avoid replacing recently used cache data and a set prediction mechanism that helps reduce L2 cache misses.

The L2 cache uses a direct mapped cache methodology. There is a dedicated external interface to the L2 cache not shared with the 6XX bus. This allows concurrent access to both the L2 cache and the 6XX bus.

The memory subsystem of the Models 6C1 and 6E1 supports a 128-bit data path to memory along with an 8-bit ECC code that provides single bit correction, double bit error detection, and 4-bit packet error detection. A minimum of 512 MB ECC SDRAM of memory is allowed, and can be expanded to a maximum of 8 GB ECC SDRAM.

The 6XX bus or system bus is optimized for high performance and multiprocessing performance. The bus is fully parity checked and each memory or cache 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 system bus speed is operated at 95.14 MHz with the 333 MHz processor card (2:7 ratio), 93.75 MHz with the 375 MHz processor card (1:4 ratio), and at 90 MHz with the 450 MHz processor card (1:5 ratio).

3.1.4 RAS features

The following are the RAS features of the p610 Model 6C1 and 6E1:

- ▶ Copper-based microprocessors
- ▶ Dynamic Processor Deallocation
- ▶ Dynamic Memory Deallocation
- ▶ ECC L2 cache and SDRAM
- ▶ Service processor
- ▶ Hot-swappable disk bays
- ▶ RAID 5 support for internal hot-swappable disks (optional)
- ▶ Optional redundant hot-plug power supplies
- ▶ Redundant hot-plug cooling fans
- ▶ Programmable service indicator lights
- ▶ Light Path Diagnostics
- ▶ Service agent
- ▶ HACMP

3.1.5 Minimum and standard features

Table 3-1 on page 37 summarizes the minimum and standard feature for the p610 Model 6C1 and 6E1.

Table 3-1 p610 Model 6C1 and 6E1 minimum and standard features

Features (min./standard)	Description
Microprocessor type	64-bit POWER3-II.
Min. number of processors	<ul style="list-style-type: none"> ▶ 1-way 333 MHz, 4 MB L2 Cache (FC 5322) ▶ 1-way 375 MHz, 4 MB L2 Cache (FC 5300) ▶ 1-way 450 MHz, 8 MB L2 Cache (FC 5309)
Memory (minimum)	512 MB.
Memory bus width	Quad 512-bit.
Data/Instruction (L1) cache	64 KB/32 KB.
Level 2 (L2) cache	4 MB of L2 cache per processor on the 333 or 375 MHz processor; 8 MB of L2 cache per processor on the 450 MHz processor.
Internal disk bays	Six available hot-swap disk drive bays with the 375 or 450 MHz processor plus two non hot-swappable disk drives. Orderable option with the 333 MHz processor.
Internal disk drive	36.4 GB Ultra SCSI.
Media bays	Four media bays: One diskette drive, one CD-ROM or DVD-RAM (DVD-RAM not available with 333 MHz processor), one non hot-swap disk drive (default), one CD-ROM, DVD-RAM, disk drive, or tape drive (optional) (DVD-RAM not available with 333 MHz processor).
PCI slots	Five PCI.
PCI bus width	32- and 64-bit.
Standard Ports	Three serial ports, one parallel, one keyboard, and one mouse ports.
Integrated SCSI Adapters	Two integrated Ultra3 SCSI controllers (one internal and one external)
Integrated LAN adapter	10/100 Mbps Ethernet controller with two ports.
Diskette drive	Yes.
System support (service) processor	Yes.
Operating system version	<ul style="list-style-type: none"> ▶ AIX 5L Versions 5.1/5.2 ▶ Linux 2.4 available from one or more IBM Linux Distribution Partners

Features (min./standard)	Description
Warranty	One year.

3.1.6 System expansion

Table 3-2 summarizes the possible maximum processor, memory, and storage expansion features for the p610 Model 6C1 and 6E1.

Table 3-2 p610 Model 6C1 and 6E1 system expansion features

System expansion features	Description
SMP configuration	<ul style="list-style-type: none"> ▶ 2-way 333 MHz, 4 MB L2 Cache (FC 5322) ▶ 2-way 375 MHz, 4 MB L2 Cache (FC 5300) ▶ 2-way 450 MHz, 8 MB L2 Cache (FC 5309)
Memory	Up to 8 GB
Internal disk drive	36.4 GB, 73.4 GB, and 146.7 GB drives available
External storage	<ul style="list-style-type: none"> ▶ IBM 2104 Expandable Storage Plus (Ultra3 SCSI) ▶ IBM 7133 Serial Disk System (SSA)

3.1.7 Features

Refer to Table 3-31 on page 150 for information about the features for the p610 Model 6C1 and 6E1.

3.1.8 Configuration notes

The following section covers the configuration note for the p610 Model 6C1 and 6E1.

Processors

The following are the configuration notes regarding the processors:

- ▶ 1-way (FC 5322) 333 MHz POWER3-II processor, 1-way (FC 5300) 375 MHz POWER3-II processor, and 1-way (FC 5309) 450 MHz POWER3-II cannot be mixed in a 2-way configuration.
- ▶ Ultra3 SCSI backplane feature number 6567 is required when ordering processor FCs 5300 or 5309. It is optional with processor FC 5322.
- ▶ Processor cards must be installed starting in slot 1 first, then slot 2.

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 8 GB of system memory.
- ▶ A minimum of 512 MB of memory is required.
- ▶ Memory DIMMs must be ordered and installed in pairs on the memory cards. They must be installed starting at the bottom of the card (card slot J1 and J2) and then moving up (Figure 3-4).
- ▶ All slots in the first memory card must be filled before a second memory card is installed.

Figure 3-4 shows the memory card locations for p610 Model 6C1 and 6E1.

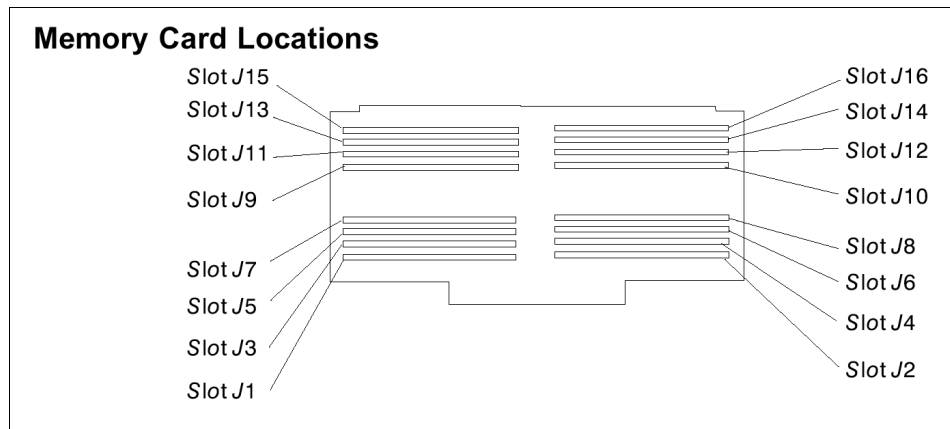


Figure 3-4 Memory card locations

Power supplies

The following are the configuration notes regarding the power supplies:

- ▶ The base machine contains two AC power supplies (FC 6277) in the lower two slots. FC 6277 is initial order only.
- ▶ AC redundant power supply (FC 6275).

Power cords

The following are the configuration notes regarding the power cords:

- ▶ Refer to Appendix C, "Power cord features" on page 907 for further information.

Racks

The following are the configuration notes regarding the racks:

- ▶ A maximum of seven pSeries 610 Model 6C1 can be mounted in a T00 rack.
- ▶ A maximum of eight pSeries 610 Model 6C1 can be mounted in a T42 rack.

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ Default DASD is in Bolt-In DASD Bay 0. Hot-swap Hard Disk Bay 1 can also be used for mandatory disk.
- ▶ FC 6156 and FC 2623 can be installed only if FC 6567 is installed.
- ▶ When FC 6156 or FC 2623 are installed in the system, the system performance may be degraded.
- ▶ Place smallest capacity DASD first, then place DASD with next highest capacity. Install DASD in ascending order capacity until all DASDs are in place.
- ▶ Media bay can accommodate a CD-ROM, DVD-RAM, tape drive, or other media device.
- ▶ All DASD and media devices are driven by the internal/integrated SCSI port only.
- ▶ Refer to Figure 3-5 on page 41 for Model 6E1, Figure 3-6 on page 42 for Model 6C1, and Table 3-3 on page 42 for additional configurations for storage devices.

Figure 3-5 on page 41 shows the disks and media bay locations for Model 6E1.

The items labeled in Figure 3-5 on page 41 are as follows:

1. Diskette drive
2. Hot-swap disk drives
3. Cover release lever
4. CD-ROM drive
5. Media bay
6. Operator panel




Figure 3-5 p610 Model 6E1 disks and media bay locations

Figure 3-6 on page 42 shows the disks and media bay locations for Model 6E1.

The items labeled in Figure 3-6 on page 42 are as follows:

1. Diskette drive
2. Hot-swap disk drives
3. Serial port connector
4. Cover release lever
5. Operator panel
6. Media bay
7. CD-ROM drive




Figure 3-6 p610 Model 6C1 disks and media bay locations

Table 3-3 describes the disks and media bay locations and the SCSI IDs.

Table 3-3 Disk and media bay locations and SCSI-IDs

Item number	Bay location	Description	SCSI-ID
1	Bay D01	Disk drive (behind operator panel)	SCSI ID 0
2	Bay D02	Media	SCSI ID 1
3	Bay D03	IDE CD-ROM (IDE Non-SCSI)	
3	Bay D03	SCSI Device	SCSI ID 2
4	Bay D04	Disk bay	SCSI ID 10
5	Bay D05	Disk bay	SCSI ID 11
6	Bay D06	Disk bay	SCSI ID 12
7	Bay D07	Disk bay	SCSI ID 13
8	Bay D08	Disk bay	SCSI ID 14
9	Bay D09	Disk bay	SCSI ID 15

Note: The SCSI bus IDs listed are the recommended values. The SCSI IDs shown for media devices indicate how the devices are set when shipped from the factory. Field installations might not comply with these recommendations.

The internal storage devices in Table 3-4 are limited to the combined maximum quantity shown below.

Table 3-4 Internal storage devices supported in systems 6C1 and 6E1

Feature code	Description	Quantity
2633	650 MB 48X CD-ROM	2
2623	4.7 GB DVD-RAM	2
6158	4 mm 20/40 GB Tape	1
6156	8 mm 20/40 GB Tape	1
3263	18.2 GB 10 K rpm Ultra3 SCSI Disk, Hot-swap	6
3102	18.2 GB 10 K rpm Ultra3 SCSI Disk, Bolt-in	2
3264	36.4 GB 10 K rpm Ultra3 SCSI Disk, Hot-swap	6
3119	36.4 GB 10 K rpm Ultra3 SCSI Disk, Bolt-in	2
3265	73.4 GB 10 K rpm Ultra3 SCSI Disk, Hot-swap	6
3118	73.4 GB 10 K rpm Ultra3 SCSI Disk, Bolt-in	2

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ Slots 1 and 2 are 32-bit and run at 33 MHz.
- ▶ Slot 3 is 64-bit and runs at 33 MHz.
- ▶ Slots 4 and 5 are 64-bit and run at 50 MHz.
- ▶ 32-bit slots are half-size.
- ▶ 64-bit slots are full-size.
- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of one per FC 6230).
- ▶ The Gigabit Fibre Channel Adapter (FC 6228) is restricted to only one adapter in slots 1 or 2 and/or one adapter in slots 3, 4, or 5.
- ▶ FC 2498 and FC 4258 are required for internal RAID. FC 4258 requires FC 2498 or FC 6203, or a follow-on adapter to be installed.
- ▶ Refer to Appendix B, “Adapter placement guidelines” on page 785 for additional I/O adapter configuration notes.

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ Graphics adapter, keyboard, and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported in the pSeries 640 is one.

Hot-plug options

The following are the configuration notes regarding the Hot-Plug options:

- ▶ It is not necessary to power down the system to install/remove/replace hot-plug options. Refer to the product documentation for further information.
- ▶ The following options are hot-pluggable:
 - Hot-plug power supplies
 - Hot-plug disk drive units (DASD)
 - Hot-plug fans

3.1.9 Express configurations

The following are the express configurations available for the Model 6C1 and 6E1. Additional, optional features may be added as desired. Except as noted, all standard configuration rules for Model 6C1 and 6E1 remain unchanged. Table 3-5 summarizes some of the express configuration options for the p610.

Table 3-5 p610 express configuration summary

Express config.	OS	CPU	Memory	DASD
100 C/E	AIX 5L	One 1-way 450 MHz POWER3-II	1024 MB	One 36.4 GB
200 C/E	AIX 5L	Two 1-way 450 MHz POWER3-II	2048 MB	Two 36.4 GB
250 C/E	AIX 5L	Two 1-way 450 MHz POWER3-II	4096 MB	Two 36.4 GB
125 C/E	AIX 5L	One 1-way 450 MHz POWER3-II	2048 MB	Six 36.4 GB with RAID adapter

pSeries 610 (100C/100E)

The 100C and 100E Express Configurations for the pSeries 610 Model 6C1 or 6E1 consists of:

- ▶ One FC 8091 (1-way 450 MHz POWER3-II processor)

- ▶ One FC 8093 (1024 MB memory DIMMs)
- ▶ One FC 3264 (36.4 GB hot-swap disk drive)
- ▶ One FC 6275 (Redundant power supply)
- ▶ For 7028-6C1: (FC 8096)
 - Minimum required: 0
 - Maximum allowed: 1 (Initial order maximum: 1)
 - OS level required: AIX 5L Version 5.1 or later
 - Initial Order/MES/Both/Supported: Initial
 - CSU: Does not apply
 - Return parts MES: Does not apply

pSeries 610 (200C/200E)

The 200C or 200E Express Configurations for the pSeries 610 Model 6C1 or 6E1 consist of:

- ▶ Two FC 8092 (1-way 450 MHz POWER3-II processor)
- ▶ Two FC 8094 (1024 MB memory DIMMs)
- ▶ Two FC 3264 (36.4 GB hot-swap disk drive)
- ▶ One FC 6275 (Redundant power supply)
- ▶ For 7028-6C1: (FC 8097)
 - Minimum required: 0
 - Maximum allowed: 1 (Initial order maximum: 1)
 - OS level required: AIX 5L Version 5.1 or later
 - Initial Order/MES/Both/Supported: Initial
 - CSU: Does not apply
 - Return parts MES: Does not apply

pSeries 610 (250C/250E)

The 250C or 250E Express Configurations for the pSeries 610 Model 6C1 or 6E1 consist of:

- ▶ Two FC 8092 (1-way 450 MHz POWER3-II processor)
- ▶ Four FC 8095 (1024 MB memory DIMMs)
- ▶ Two FC 3265 (73.4 GB hot-swap disk drive)
- ▶ One FC 6275 (Redundant power supply)

- ▶ For 7028-6C1: (FC 8098)
 - Minimum required: 0
 - Maximum allowed: 1 (Initial order maximum: 1)
 - OS level required: AIX 5L Version 5.1 or later
 - Initial Order/MES/Both/Supported: Initial
 - CSU: Does not apply
 - Return parts MES: Does not apply

pSeries 610 (125C/125E) RAID applications

The 125C or 125E Express Configurations for the pSeries 610 Model 6C1 or 6E1 consist of:

- ▶ One FC 8091 (1-way 450 MHz POWER3-II processor)
- ▶ Two FC 8094 (1024 MB memory DIMMs)
- ▶ One FC 3119 (36.4 GB disk drive)
- ▶ Six FC 3264 (36.4 GB hot-swap disk drive)
- ▶ One FC 6275 (Redundant power supply)
- ▶ One FC 2498 (Ultra3 SCSI RAID Adapter)
- ▶ For 7028-6C1: (FC 8099)
 - Minimum required: 0
 - Maximum allowed: 1 (Initial order maximum: 1)
 - OS level required: AIX 5L Version 5.1 or later
 - Initial Order/MES/Both/Supported: Initial
 - CSU: Does not apply
 - Return parts MES: Does not apply

3.2 7029 Model 6C3 and 6E3 IBM @server pSeries 615

The IBM @server® pSeries 615 Models 6C3 and 6E3 (referred to hereafter as the Model 6C3 and Model 6E3, or p615 when discussing both models) are designed for customers looking for a cost-effective, high performance, space-efficient server that uses advanced IBM technology first available in the high-end pSeries 690, then on the midrange product p650, and finally in the entry-level p630. The p615 uses the 64-bit, copper and SOI-based POWER4+ microprocessor, and is intended for use in stand-alone or LAN clustered environments.

The p615 is a member of the 64-bit family of symmetric multiprocessing (SMP) UNIX servers from IBM. The Model 6C3 (product number 7029-6C3) is a 4-EIA (4U), 19-in. rack-mounted server, while the Model 6E3 (product number 7029-6E3) is a desktop server. The p615 can be configured as a 1- or 2-way system running at 1.2 GHz or as a 2-way running at 1.45 GHz. At the time of writing, the total system memory can range from 1 GB to 16 GB based on the available DIMMs. The availability of cluster support enhances the already exceptional value of these models.

The Model 6C3 rack-mounted server is intended to be installed in a 19-in. rack, thereby enabling efficient use of computer room floor space. The Model 6C3 is designated as a customer setup system and requires three persons (due to weight and safety issues) to be present to install the unit into the rack.

The Model 6E3 is a desktop server, ideal for environments requiring the user to have local access to the machine. A typical example of this would be applications requiring a native graphics display.

The GXT135P 2D graphics accelerator with analog and digital interfaces (FC 2849) is available and is supported for SMS, firmware menus, and other low-level functions, as well as when AIX starts the X11-based graphical user interface. Graphical AIX system tools are usable for configuration management if the adapter is connected to the primary console, such as the IBM L200p flat-panel monitor (FC 3636) or the IBM T541H 15-in. TFT Color Monitor (FC 3637).

Each system is delivered preconfigured as ordered. In the case of the p615 Linux ready Express Configurations, the systems are delivered without an operating system. You can order SuSE SLES 8 at the time of purchase. The order will be fulfilled by SuSE. The system is designed to be set up by the customer and, in most cases, will not require the use of any tools. Full setup instructions are included with the system.

Figure 3-7 on page 48 shows the pSeries 615 Model 6C3 and 6E3 physical packaging.



Figure 3-7 p615 Model 6E3 and 6C3 physical packaging

3.2.1 Product positioning

With their performance, and SMP capability, the Model 6E3 and Model 6C3 are well-positioned for transaction-intensive server solutions such as e-commerce, database, and application database server applications.

The Model 6E3 and Model 6C3 provide 2-way SMP capability, up to 16 GB of memory, and up to 1174.4 GB of internal DASD, which make an excellent solution for many demanding commercial applications. It is a cost-effective performance alternative to the IBM @server pSeries p610 Models 6E1 and 6C1.

The Model 6E3 and Model 6C3 also offer a variety of high availability features, such as:

- ▶ An integrated service processor
- ▶ Hot-swap redundant power supply (AC)
- ▶ Hot-swap cooling fans and disk drives
- ▶ LED Service Identification
- ▶ Auto reboot on power loss

Figure 3-8 on page 49 shows the rack view of the Model 6C3 and the workstation view of the Model 6E3.



Figure 3-8 p615 Model 6C3 rack view and Model 6E3 desktide view

3.2.2 Highlights

The pSeries 615 Model 6C3 and 6E3 deliver a cost-efficient growth path to the future through such attributes as:

- ▶ Outstanding value for entry UNIX servers
- ▶ Powerful POWER4+ processor technology, symmetric multiprocessing (SMP), and e-business server
- ▶ Rack-mount (Model 6C3) configuration
- ▶ Desktide (Model 6E3) configuration
- ▶ 64-bit scalability in 1- and 2-way configurations with the following processors options:
 - 1- or 2-way 1.2 GHz POWER4+ with 8 MB of L3 cache
 - 2-way 1.45 GHz POWER4+ with 8 MB of L3 cache
- ▶ Up to 16 GB of ECC SDRAM memory
- ▶ Up to 1174.4 GB of internal disk storage

- ▶ Two slimline media bays for DVD-ROM and diskette
- ▶ One media bay for tape drive or DVD-RAM
- ▶ Up to eight disk bays (four basic bays and four option disk bays)
- ▶ Six hot-swap, 64-bit PCI-X slots
- ▶ An integrated service processor
- ▶ Optional redundant hot-swap AC power
- ▶ Support for SCSI, SSA, and Fibre Channel attached storage systems
- ▶ Support for 32-bit and 64-bit applications
- ▶ Fourteen new express configurations
- ▶ Cluster 1600 enhancements: p615 CSM support for Model 6C3

3.2.3 Technical overview

This section provides the technical overview and the system architecture of the p615 Model 6C3 and 6E3 represented by Figure 3-9 on page 51. The major components of this diagram will be described in the following sections. The bandwidths provided throughout this section are theoretical maximums provided for reference. It is always recommended to obtain real-world performance measurements using production workloads.

Figure 3-9 on page 51 shows the system diagram of p615 Model 6C3 and 6E3.




Figure 3-9 Conceptual diagram of the p615 POWER4+ system architecture

The p615 implements the POWER4+ processor chip running at 1.2 GHz or 1.45 GHz. However, unlike the Multichip Module (MCM) packaging first used in the pSeries 690, or the Single Chip Module (SCM) used in pSeries 650 and 630, the p615 system uses a one- or two-core SCM that is packaged with another chip, which combines an 8 MB L3 cache with a memory controller chip that is directly interfaced to the memory buffer SMI chips as shown in the enlarged view of the processor subsystem (Figure 3-24) and in Figure 3-22. The chip is permanently mounted on the system planar, containing either one or two processor cores (CPUs). Only a one- or two-processor configuration is available.

The chip-to-chip fabric bus (which was used between chips on the same MCM in the Model 690) is no longer relevant here. This module does not have an external fabric bus.

Figure 3-10 on page 52 shows the logical diagram of the p615 processor/cache module.

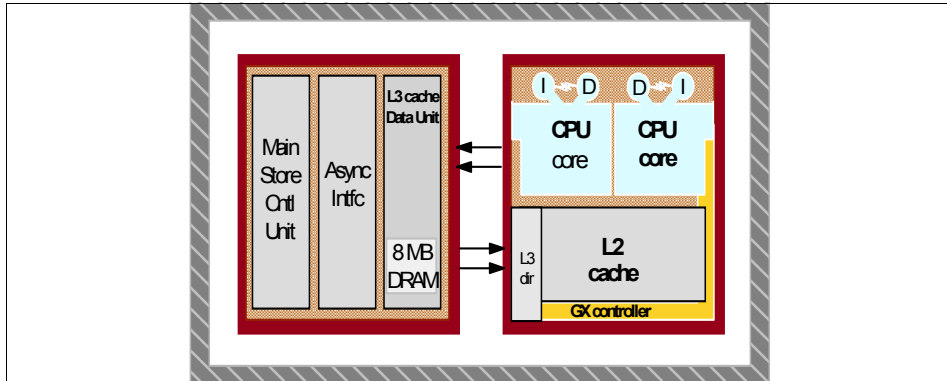


Figure 3-10 The p615 processor/cache module

Each processor/cache module has a Ceramic Column Grid Array (CCGA) package where the chip carrier is raised slightly from its board mounting by small metal solder columns that provide the required connections and improved thermal resilience characteristics, and a Land Grid Array (LGA) socket for bringup, as shown in Figure 3-11.




Figure 3-11 Enlarged view of the processor subsystem on system planar

The POWER4+ storage subsystem consists of three levels of cache and the memory subsystem. The first two levels of cache are onboard the POWER4+ chip. The first level is 64 KB of Instruction (I) cache and 32 KB of Data (D) cache per processor core. The second level is 1.5 MB of L2 cache on the POWER4+. A two-way processor shares the L2 cache. All caches have either full ECC or parity protection on the data arrays, and the L1 cache has the ability to refetch data from the L2 cache in the event of soft errors detected by parity checking.

The Level 3 cache consists of two components: the L3 cache controller/directory and the L3 data array. The L3 cache controller/directory is on the POWER4+ chip, and the L3 data array, which consists of 8 MB of embedded DRAM (eDRAM), is located in the other chip of the processor/cache module. In 2-way systems, the L3 cache is shared.

The conceptual diagram of the memory subsystem of the p615 is shown in Figure 3-12. There are four 8-byte data paths from the memory controller to the memory with an aggregated bandwidth of 6.4 GBps to the processor card.

DDR memory can theoretically double memory throughput at a given clock speed by providing output on both the rising and falling edges of the clock signal (rather than only on the rising edge).

Memory access is through the on-chip L2 cache and Level 3 cache to the off-chip synchronous memory interface (SMI) and finally to the memory DIMMs, as represented in Figure 3-12.

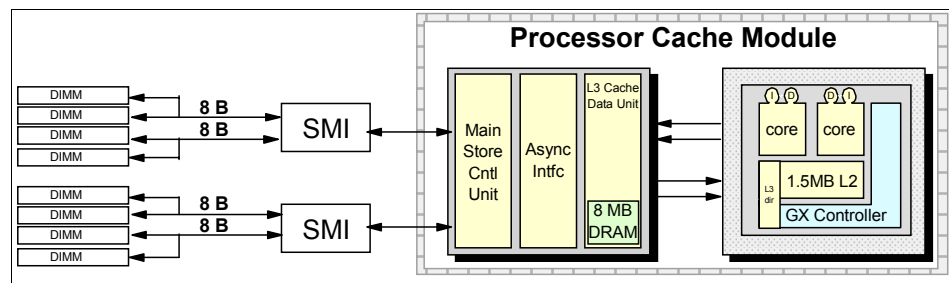


Figure 3-12 Conceptual diagram of POWER4+ processor and memory subsystem

The p615 uses Error Checking and Correcting (ECC) circuitry for system memory to correct single-bit and to detect double-bit memory failures. Detection of double-bit memory failures helps maintain data integrity. Furthermore, the memory chips are organized such that the failure of any specific memory module only affects a single bit within a four-bit ECC word (*bit-scattering*), thus allowing for error correction and continued operation in the presence of a complete chip failure (*Chipkill recovery*). The memory DIMMs also utilize *memory scrubbing* and thresholding to determine when spare memory modules within each bank of memory should be used to replace ones that have exceeded their threshold of error count (*dynamic bit-steering*). Memory scrubbing is the process of reading the contents of the memory during idle time and checking and correcting any single-bit errors that have accumulated by passing the data through the ECC logic. This function is a hardware function on the memory controller chip and does not influence normal system memory performance. See Figure 3-13 on page 54 for reference.

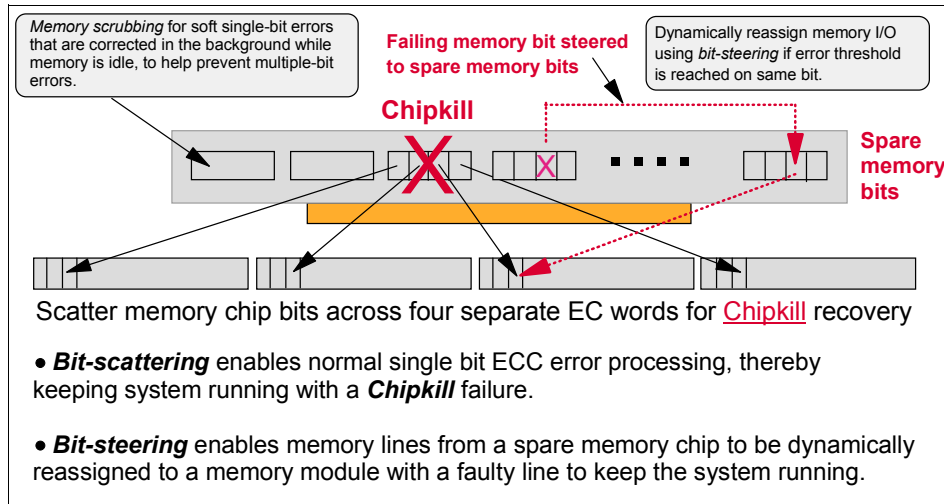


Figure 3-13 Main storage ECC and extensions

The system bus from the processors to the memory subsystem is 2 x 16 bytes at one third the CPU clock speed for an aggregate data rate of 12.8 GBps. All system, I/O, and PCI-X buses support parity error detection. The memory controller is integrated on the processor/cache module and directly interfaced with two 8-byte memory data ports (SMI) to control system DRAM.

The GX controller (embedded in the POWER4+ chip) is responsible for controlling the flow of data through the GX bus. The GX bus is a high-frequency, single-ended, unidirectional, point-to-point bus. Both data and address information are multiplexed onto the bus, and for each path there is an identical bus for the return path. The GX bus has dual 4-byte paths at 400 MHz to give an aggregate data rate of 3.2 GBps (for a 1.2 GHz CPU). The processor card connects to the GX bus through its GX controller.

The I/O bridge contains a series of 1-byte buses grouped into pairs called ports. One of the buses in the pair is for inbound data transfers (500 MBps), while the other bus is for outbound data transfers (500 MBps). There is a total of two ports for the internal I/O. The two internal ports connect to a PCI-X Host Bridge (PHB). The PHB chip acts as a bridge between the I/O bus and two PCI-X to PCI-X bridges, which fan out to integrated I/O controllers and slots.

The PCI-X slots meet 64-bit 133 MHz requirements. However, two slots (5 and 6) are affected in performance by the EADS-X chip 1 that is connected to primary CSP interface. This lowers the speed of the EADS-X and its connected PCI-X Host Bridge port to 66 MHz. PCI-X and PCI adapters that demand high performance should be installed into slots 1 and 4 driven by the EADS-X chip 2

that is not connected to the CSP interface. Slots 2 and 3 are still connected to the EADS-X chip 2, but they are 32-bit slots.

PCI-X is the latest version of PCI bus technology, using a higher clock speed (133 MHz) to deliver a bandwidth of up to 1 GBps. The PCI-X slots in the p615 system support hot-plug and Extended Error Handling (EEH). EEH-capable adapters respond to a specially generated data packet from a PCI-X slot with a problem. This packet is analyzed by the system firmware, which then allows the device driver to reset the adapter or slot, isolating the error and reducing the need for a system reboot.

The PCI-X slots support existing 3.3 volt PCI adapters, which enables them to support both 64-bit and 32-bit adapters.

Choosing between 32-bit and 64-bit adapters influences slot placement and affects performance. Higher-speed adapters use 64-bit slots because they can transfer 64 bits of data for each data transfer phase.

Generally, 32-bit adapters can function in 64-bit PCI-X slots; however, 64-bit adapters cannot be used in the 32-bit slots on p615 systems.

3.2.4 RAS features

The following are the RAS features of the p615 Model 6C3 and 6E3:

- ▶ Fault avoidance through highly reliable component selection, component minimization, and error handling technology designed into the chips
- ▶ Improved reliability through processor operation at a lower voltage, enabled by the use of copper chip circuitry and Silicon-on-Insulator technology
- ▶ Fault tolerance through additional hot-swappable power supply, and the capability to perform concurrent maintenance for power and cooling
- ▶ Automatic First Failure Data Capture (FFDC) and diagnostic fault isolation capabilities
- ▶ Concurrent run-time diagnostics based on First Failure Data Capture
- ▶ Predictive failure analysis on processors, cache, memory, and disk drives
- ▶ Dynamic error recovery
- ▶ Error Checking and Correction (ECC) or equivalent protection (such as refetch) on main storage, all cache levels (1, 2, and 3), and internal processor arrays
- ▶ Dynamic processor deallocation based on run-time errors (requiring more than one processor)

- ▶ Persistent processor deallocation (boot-time deallocation based on run-time errors)
- ▶ Persistent deallocation extended to memory
- ▶ Chipkill correction in memory
- ▶ Memory scrubbing and redundant bit-steering for self-healing
- ▶ Industry-leading PCI-X bus parity error recovery as first introduced on the p690 systems
- ▶ Hot-plug functionality of the PCI-X bus I/O subsystem
- ▶ PCI-X bus and slot deallocation
- ▶ Disk drive fault tracking that monitors the number/rate of data errors and thresholds several recoverable hardware errors
- ▶ Avoiding checkstops with process error containment
- ▶ Environmental monitoring (temperature and power supply)
- ▶ Auto-reboot

3.2.5 Minimum and standard features

This section describes the system minimum and standard features for p615 Model 6C4 and 6E4.

Table 3-6 summarizes the minimum and standard features for p615 Model 6C3 and 6E3.

Table 3-6 p615 Model 6C3 and 6C4 minimum and standard features

Features (min./standard)	Description
Microprocessor type	POWER4+.
Min. number of processors	1-way POWER4+ 1.2 GHz and 8 MB L3 Cache (FC 5220).
Memory (minimum)	1 GB.
Data/Instruction (L1) cache	64 KB/32 KB.
Level 2 (L2) cache	1.5 MB.
Internal disk bays	Four hot-swap disk drive bays.

Features (min./standard)	Description
Internal disk drive	Each bay can contain one of the following disks: <ul style="list-style-type: none"> ▶ 36.4 GB Ultra320 10K RPM (FC 3273) ▶ 36.4 GB Ultra320 15K RPM (FC 3277) ▶ 73.4 GB Ultra320 10K RPM (FC 3274) ▶ 73.4 GB Ultra320 15K RPM (FC 3278) ▶ 146.8 GB Ultra320 10K RPM (FC 3275)
Media bays	Three media bays: <ul style="list-style-type: none"> ▶ One DVD-ROM ▶ One diskette drive (optional) ▶ One DVD-RAM or tape drive (optional)
PCI slots	Six PCI-X slots with 1.2 or 1.45 GHz processor systems. Hot-swap, 64-bit, 133 MHz, and 3.3V.
Standard ports	Three serial ports (front serial port available on Model 6C3 only, one parallel, one keyboard, and one mouse ports).
Integrated SCSI adapters	Two integrated Ultra3 SCSI controllers (internal only).
Integrated LAN adapter	10/100 Mbps Ethernet controller with two ports.
Fans	Hot-swap and redundant fans.
Power	One hot-swap power supply (Redundant power optional).
System support (service) processor	Yes.
Operating system version	AIX 5L Version 5.1 or 5.2, or later, operating system or SuSE Linux Enterprise Server 8 for iSeries and pSeries systems.
Warranty	One year.

3.2.6 System expansion

This section describes the available features for the p615 Model 6C4 and 6E4 expansion.

Table 3-7 on page 58 summarizes the possible maximum processor, memory, and storage expansion features for p615 Model 6C3 and 6E3.

Table 3-7 p615 Model 6C4 and 6E4 system expansion features

System expansion features	Description
SMP configuration	<ul style="list-style-type: none"> ▶ 2-way POWER4+ 1.2 GHz, 8 MB L3 cache. ▶ 2-way POWER4+ 1.45 GHz, 8 MB L3 cache.
Memory	<p>Up to 16 GB using the following available order combinations:</p> <ul style="list-style-type: none"> ▶ FC 4446, 1024 MB (4 x 256 MB) DIMMs, 208-pin 8 ns DDR SDRAM. ▶ FC 4447, 2048 MB (4 x 512 MB) DIMMs, 208-pin 8 ns DDR SDRAM. ▶ FC 4448, 4096 MB (4 x 1024 MB) DIMMs, 208-pin 8 ns DDR SDRAM. ▶ FC 4449, 8192 MB (4 x 2048 MB) DIMMs, 208-pin 8 ns DDR SDRAM.
Disk drive and disk bay expansion	<ul style="list-style-type: none"> ▶ Up to 1.17 TB (36.4 GB, 73.4 GB, and 146.8 GB drives available). ▶ Up to eight bays.
External Storage	<ul style="list-style-type: none"> ▶ IBM 2104 Expandable Storage Plus (Ultra3 SCSI). ▶ IBM 7133 Serial Disk System (SSA).

3.2.7 Features

Refer to Table 3-31 on page 150 for information about the features for p615 Model 6C3 and 6E3.

3.2.8 Configuration notes

The following section covers the configuration notes for the p615 Model 6C3 and 6E3.

Figure 3-14 on page 59 shows a detailed view of the Model 6C3 and Model 6E3, including the locations of all various components and devices.




Figure 3-14 Models 6C4 and 6E4 (POWER4+ system with six PCI-X slots)

Processors

The following are the configuration notes regarding the processors:

- ▶ Only one processor card can be installed in the system.
- ▶ 1- or 2-way SMP system
 - 1- or 2-way, 64-bit, copper-based POWER4+ microprocessors running at 1.2 GHz, with a total of 8 MB L3 cache.
 - 1-way system planar (FC 5220, or Express Configuration FC 8148)
 - 2-way system planar (FC 5222, or Express Configuration FC 8149)
 - 2-way, 64-bit, copper-based POWER4+ microprocessor running at 1.45 GHz with 8 MB of L3 cache shared between the CPUs (FC 5234, or Express Configuration FC 8187)

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 16 GB of system memory.
- ▶ A minimum of 1 GB of memory is required.
- ▶ Memory DIMMs must be ordered and installed in quads on the processor card. They must be installed starting at the bottom of the card.
 - First quad:
 - J0A, J1A, J2A, J3A
 - Second quad:
 - J0B, J1B, J2B, J3B

Figure 3-15 shows the memory DIMM slot locations associated with each memory DIMM slot.

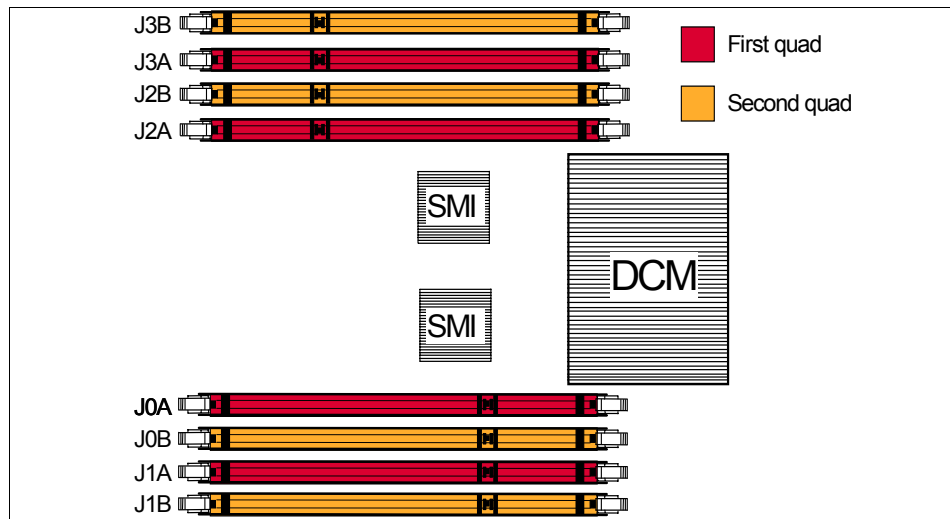


Figure 3-15 Memory placement for the p615

- ▶ The p615 does not officially support any OEM memory, and there is no exception to this rule. OEM memory is never certified for use in pSeries servers. If the p615 is populated with OEM memory, you could experience unexpected and unpredictable behavior.
- ▶ All IBM memory is identified by an IBM logo and a white label printed with a barcode on top and an alphanumeric string on the bottom, created according to the rule reported in Figure 3-16 on page 61.




Figure 3-16 IBM barcode label

Power supplies

The following are the configuration notes regarding the power supplies:

- ▶ Base machine contains one AC power supply with a second available for redundancy
- ▶ Second optional AC (FC 6266) power supply

Power cords

The following are the configuration notes regarding the power cords:

- ▶ Refer to Appendix C, “Power cord features” on page 907 for further information.

Racks

The following are the configuration notes regarding the racks:

- ▶ The Model 6C3 is designed to be placed at any location in the rack. For rack stability, it is advisable to start filling a rack from the bottom.
- ▶ A maximum of nine pSeries 615 Model 6C4s can be mounted in a T00 rack.
- ▶ A maximum of 10 pSeries 615 Model 6C4s can be mounted in a T42 rack.
- ▶ Any remaining space in the rack can be used to install other systems or peripherals, provided that the maximum permissible weight of the rack is not exceeded and the installation rules for these devices are followed.
- ▶ The Model 6C3 can be installed in a suitable OEM rack, provided that the rack conforms to the EIA-310-D standard. This standard is published by the Electrical Industries Alliance, and a summary of this standard is available in the publication *RS/6000 and IBM @server pSeriesSite and Hardware Planning Information*, SA38-0508. An online copy of this document can be found at:

http://www.ibm.com/servers/eserver/pseries/library/hardware_docs

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ There are two DASD 4-packs in the system, one mandatory (FC 6574) and one optional (FC 6574).
- ▶ Media bays can contain an optional DVD-ROM in bay 3, an optional diskette in bay 2, and an optional DVD-RAM or tape drive in bay 4.
- ▶ DVD-RAM (FC 2623) requires SCSI adapter (FC 6203 or follow-on) and cable (FC 4266).
- ▶ If optional tape drive FC 6158, FC 6120, or FC 6134 is installed, cable FC 4263 is required. Cable FC 4263 requires FC 6158, FC 6120, or FC 6134 to be installed.
- ▶ Place smallest-capacity DASD first, then place DASD with next highest capacity. Install DASD in ascending order capacity until all DASDs are in place.
- ▶ Boot DASD is placed in DASD #8. If the system has one 4-pack installed, DASD is filled #7, #6, #5. If two 4-packs are installed, DASD is filled #4, #7, #3, #6, #2, #5, #1.
- ▶ Boot is supported via:
 - CD-ROM/DVD-ROM/DVD-RAM
 - Internal or external tape drives
 - SCSI disk
 - SSA disk
 - SAN boot using a 2 GB Fibre Channel Adapter FC 6239
 - LAN boot

Figure 3-17 on page 63 shows the p615 bay locations.




Figure 3-17 p615 bay locations

Table 3-8 describes the operator panel, diskette, disks, and media bay locations, and SCSI IDs, if applicable.

Table 3-8 Disk and media bay locations and SCSI IDs

Item number	Bay location	Description	SCSI ID
1	Bay D01	Operator Panel	
2	Bay D02	Diskette Drive or IDE CD-ROM	
3	Bay D03	IDE CD-ROM (Default) or IDE DVD-ROM (Optional)	IDE (Non-SCSI)
4	Bay D04	Tape Drive (Optional)	SCSI ID 0
		SCSI DVD-RAM (Optional)	SCSI ID 1
5	Bay D10	Disk bay	SCSI ID 8
6	Bay D11	Disk bay	SCSI ID 9
7	Bay D12	Disk bay	SCSI ID 10
8	Bay D13	Disk bay	SCSI ID 11

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ Slots 4, 5, and 6 are long, 64-bit, 3.3V, and run at 133 MHz.
- ▶ Slots 2 and 3 are short 32-bit, 3.3V, and run at 66 MHz.
- ▶ Slot 1 is long, 64-bit, and runs at 133 MHz. Selected cards only are allowed in this slot.

Figure 3-18 shows the PCI-X locations of Model 6C3 as example for both systems.




Figure 3-18 p615 Model 6C3 PCI-X slot locations

- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of one per FC 6230).
- ▶ 64-bit adapters are not allowed in 32-bit slots.
- ▶ Refer to Appendix B, “Adapter placement guidelines” on page 785 for additional I/O adapter configuration notes.

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ Graphics adapter, keyboard, and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported is four.

Hot-plug options

The following are the configuration notes regarding the Hot-Plug options:

- ▶ It is not necessary to power down the system to install/remove/replace hot-plug options. Refer to the p615 product documentation for further information.
- ▶ The following options are hot-pluggable:
 - Hot-plug power supplies
 - Hot-plug disk drive units (DASD)
 - Hot-plug fans
 - Hot-plug adapters

3.2.9 Express configurations

The following are the express configurations available for the Model 6C3 and 6E3. Additional, optional features may be added as desired. Except as noted, all standard configuration rules for Model 6C3 and 6E3 remain unchanged. Table 3-9 provides a summary of the p615 express configurations.

Table 3-9 p615 express configuration summary

Express config.	OS	CPU	Memory	DASD
110 C/E	AIX 5L	One 1-way 1.2 GHz POWER4+	1024 MB	One 36.4 GB
111 C/E	Linux	One 1-way 1.2 GHz POWER4+	1024 MB	One 36.4 GB
120 C/E	AIX 5L	One 1-way 1.2 GHz POWER4+	2048 MB	One 36.4 GB
121 C/E	Linux	One 1-way 1.2 GHz POWER4+	2048 MB	One 36.4 GB
140 C/E	AIX 5L	One 1-way 1.2 GHz POWER4+	4096 MB	One 36.4 GB
141 C/E	Linux	One 1-way 1.2 GHz POWER4+	4096 MB	One 36.4 GB
180 C/E	AIX 5L	One 1-way 1.2 GHz POWER4+	8192 MB	One 36.4 GB
181 C/E	Linux	One 1-way 1.2 GHz POWER4+	8192 MB	One 36.4 GB
220 C/E	AIX 5L	One 2-way 1.2 GHz POWER4+	2048 MB	Two 36.4 GB
221 C/E	Linux	One 2-way 1.2 GHz POWER4+	2048 MB	Two 36.4 GB
240 C/E	AIX 5L	One 2-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
241 C/E	Linux	One 2-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
280 C/E	AIX 5L	One 2-way 1.2 GHz POWER4+	8192 MB	Two 36.4 GB

Express config.	OS	CPU	Memory	DASD
281 C/E	Linux	One 2-way 1.2 GHz POWER4+	8192 MB	Two 36.4 GB
224 C/E	AIX 5L	One 2-way 1.45 GHz POWER4+	2048 MB	Two 36.4 GB
225 C/E	Linux	One 2-way 1.45 GHz POWER4+	2048 MB	Two 36.4 GB
244 C/E	AIX 5L	One 2-way 1.45 GHz POWER4+	4096 MB	Two 36.4 GB
245 C/E	Linux	One 2-way 1.45 GHz POWER4+	4096 MB	Two 36.4 GB
284 C/E	AIX 5L	One 2-way 1.45 GHz POWER4+	8192 MB	Two 36.4 GB
285 C/E	Linux	One 2-way 1.45 GHz POWER4+	8192 MB	Two 36.4 GB

110C for Model 6C3, 110E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 110C or 110E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (110C or 110E) (FC 8167)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 1024 MB Memory DIMMS (FC 8156)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 111C for Model 6C3, 111E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 111C or 111E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (111C or 111E) (FC 8160)

- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 1024 MB Memory DIMMS (FC 8156)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

120C for Model 6C3, 120E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 120C or 120E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (120C or 120E) (FC 8169)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 2048 MB Memory DIMMS (FC 8157)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 121C for Model 6C3, 121E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 121C or 121E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (121C or 121E) (FC 8161)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 2048 MB Memory DIMMS (FC 8157)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

140C for Model 6C3, 140E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 140C or 140E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (140C or 140E) (FC 8170)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other

than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 141C for Model 6C3, 141E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 141C or 141E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (141C or 141E) (FC 8162)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

180C for Model 6C3, 180E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 180C or 180E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (180C or 180E) (FC 8171)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 8192 MB Memory DIMMS (FC 8159)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 181C for Model 6C3, 181E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 181C or 181E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (181C or 181E) (FC 8163)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8148)
- ▶ One 8192 MB Memory DIMMS (FC 8159)
- ▶ One 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

220C for Model 6C3, 220E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 220C or 220E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (220C or 220E) (FC 8172)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 2048 MB Memory DIMMS (FC 8157)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)

- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 221C for Model 6C3, 221E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 221C or 221E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (221C or 221E) (FC 8164)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 2048 MB Memory DIMMS (FC 8157)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

240C for Model 6C3, 240E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 240C or 240E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (240C or 240E) (FC 8173)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)

- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The FC 8xxx are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 241C for Model 6C3, 241E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 241C or 241E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (241C or 241E) (FC 8165)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

280C for Model 6C3, 280E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 280C or 280E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (280C or 280E) (FC 8175)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 8192 MB Memory DIMMS (FC 8159)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)

- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 281C for Model 6C3, 281E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 281C or 281E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (281C or 281E) (FC 8166)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8149)
- ▶ One 8192 MB Memory DIMMS (FC 8159)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

224C for Model 6C3, 224E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 224C or 224E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (280C or 280E) (FC 8188)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 2048 MB Memory DIMMS (FC 8157)

- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 225C for Model 6C3, 225E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 225C or 225E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (281C or 281E) (FC 8191)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 2048 MB Memory DIMMS (FC 8157)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

244C for Model 6C3, 244E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 244C or 244E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (280C or 280E) (FC 8189)

- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 245C for Model 6C3, 245E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 245C or 245E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (281C or 281E) (FC 8192)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 4096 MB Memory DIMMS (FC 8158)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

284C for Model 6C3, 284E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 284C or 284E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (280C or 280E) (FC 8189)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 8192 MB Memory DIMMS (FC 8192)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux-Ready 285C for Model 6C3, 285E for Model 6E3

One Model 6C3 or 6E3 is part of Express Configuration 285C or 285E, including a minimum of:

- ▶ One feature number indicator for Express Configuration p615 (281C or 281E) (FC 8192)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8187)
- ▶ One 8192 MB Memory DIMMS (FC 8192)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3273)
- ▶ One 4-pack disk enclosure (FC 6574)
- ▶ One diskette (FC 2611)
- ▶ One AC power supply (FC 6266)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C3 only)/FC 98xx)

The 8xxx feature numbers are mandatory. Additional, optional 7029-6C3/6E3 features may be added as desired. Upgrades may be made to the minimum

configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

3.3 7028 Model 6C4 and 6E4 IBM @server pSeries 630

The IBM @server pSeries 630 Models 6C4 and 6E4 (referred to hereafter as the Model 6C4 and Model 6E4 or p630 when discussing both models) are designed for customers looking for cost-effective, high performance, space-efficient servers that make use of IBM technology first used in the high-end pSeries 690. These systems use the 64-bit, copper/SOI-based, POWER4 and POWER4+ microprocessors, packaged as 1- and 2-way cards.

The Models 6C4 and 6E4 are members of the 64-bit family of symmetric multiprocessing (SMP) UNIX servers from IBM. The Model 6C4 (product number 7028-6C4) is a 4 EIA (4U 19-in.) rack-mounted server, while the Model 6E4 (product number 7028-6E4) is a desktide/desktop server. With a maximum of two processor cards, the Models 6C4 and 6E4 can be configured into 1-, 2-, or 4-way systems. Each processor card is packaged together with up to 16 GB of memory per card into a processor book (a sealed unit that protects the components in a rigid structure designed for higher reliability). Total system memory can range from 1 GB up to 32 GB on a 4-way system based on the currently available DIMMs.

The availability of dynamic logical partitioning (LPAR) and cluster support enhances the already exceptional value of these models. Figure 3-19 on page 78 shows the pSeries 630 Model 6C4 and 6E4.

For more information on LPAR, see 16.1, "Partitioning" on page 686.



Figure 3-19 p630 Model 6C4 and 6E4

3.3.1 Product positioning

The IBM @server pSeries 630 Models 6E4 and 6C4 are UNIX servers for customers looking for a cost-effective, high-performance UNIX server. These are customers who are attracted to the p650 but do not require the overall performance or scaling.

The Model 6C4 is an ideal replacement for the current successful Model B80. The Model 6E4 is an ideal replacement for the Model 270. Both systems use the POWER4 or POWER4+ technology, have four hot-swap PCI-X slots with the 1.0 GHz processor or six hot-swap PCI-X slots with the 1.2 or 1.45 GHz processors, four front accessible hot-swap disk bays, and can accommodate up to 32 GB of ECC memory. Two media bays are available for a CD-ROM, DVD-RAM, diskette, or tape drive. Integrated features include two 10/100 Ethernet ports, two Ultra3 SCSI (internal and external) ports, three serial, one parallel, two HMC, keyboard, and mouse ports.

Reliability improvements help make the Models 6E4 and 6C4 very competitive. Its CPU deallocation function 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 system administrator or service personnel of the condition. The computing operation continues with the processor deactivated, allowing repair to be scheduled at a time convenient to your processing schedule. Hot-plug functionality of the PCI I/O subsystem is designed to bring

new levels of up-time to your system for the replacement or addition of I/O adapters. Selective disablement of an adapter can be made as necessary, while system operations not dependent on this adapter continue to run without being affected.

High availability continues to be an ongoing and growing need for customers. IBM continues its tradition of providing high-availability solutions with the Models 6E4 and 6C4, which can be ordered in a high-availability configuration with a choice of components.

Figure 3-20 shows the rack view of the Model 6C4 and the workstation view of the Model 6E4.



Figure 3-20 p630 Model 6C4 rack view, and Model 6E4 workstation view

3.3.2 Highlights

The pSeries 630 Model 6C4 and 6E4 deliver a cost-efficient growth path to the future through such attributes as:

- ▶ POWER4 and POWER4+ processor technology

- ▶ Rack-mount (Model 6C4) configuration
- ▶ Desktop (Model 6E4) configuration
- ▶ 64-bit scalability in 1-, 2-, and 4-way configurations with the following processors options:
 - 1- or 2-way 1.2 GHz POWER4+ with 8 MB of L3 cache per processor card
 - 1- or 2-way 1.45 GHz POWER4+ with 8 MB of L3 cache per processor card
- ▶ Up to 32 GB of memory
- ▶ Up to 587.2 GB of internal disk storage
- ▶ Two media bays and four disk bays
- ▶ Six hot-swap, 64-bit PCI-X slots with 1.2 and 1.45 GHz processor systems
- ▶ An integrated service processor
- ▶ Optional hot-swap fans
- ▶ Optional redundant hot-swap AC power
- ▶ Support for SCSI, SSA, and Fibre Channel attached storage systems
- ▶ Support for 32-bit and 64-bit applications
- ▶ AIX license included

3.3.3 Technical overview

This section provides the technical overview and the system architecture of the p630 Model 6C4 and 6E4 represented by Figure 3-21 on page 81. The major components of this diagram will be described in the following sections. The bandwidths provided throughout this section are theoretical maximums provided for reference. It is always recommended to obtain real-world performance measurements using production workloads.

Figure 3-21 on page 81 shows the system diagram of p630 Model 6C4 and 6E4.

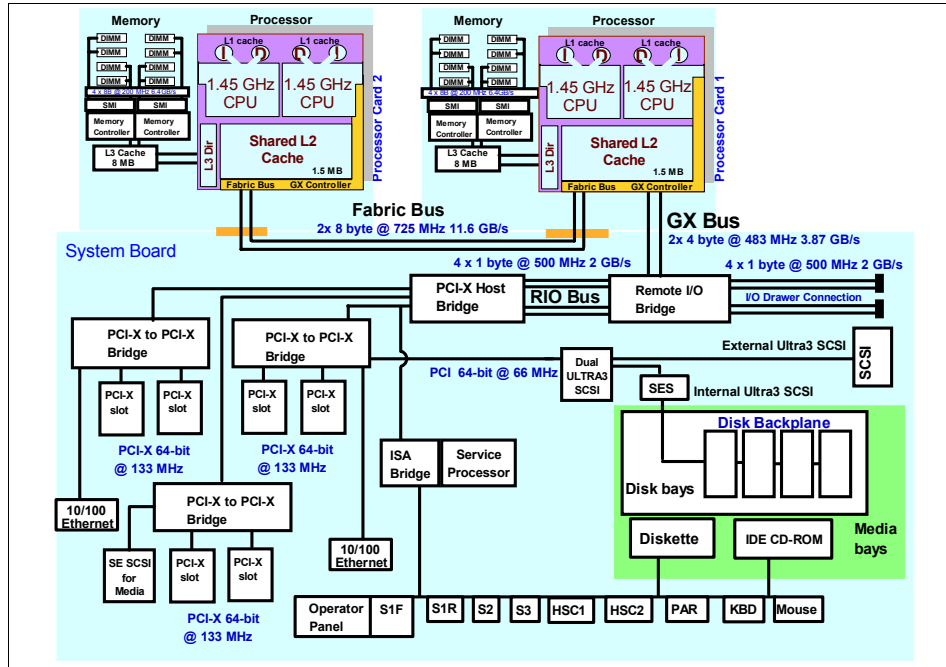


Figure 3-21 Diagram of Models 6C4 and 6E4 POWER4+ system architecture

The Model 6C4 and Model 6E4 subsystems consist of POWER4+ processors running at 1.2 GHz or 1.45 GHz. These are related to the processors used in the IBM @server pSeries 690 and 650. However, unlike the Multichip Module (MCM) packaging used in the Model 690, these systems use a Single Chip Module (SCM) containing either one or two processor cores (CPUs), with each SCM permanently mounted on a processor card. Models 6C4 and 6E4 can contain one or two processor cards, giving the option of one, two, or four CPUs per system.

One key difference is that the chip-to-chip fabric bus (which was used between chips on the same MCM in the Model 690) is no longer relevant in the SCM. It is replaced by a module-to-module fabric.

Figure 3-22 on page 82 shows the comparison between the SCM and MCM modules.




Figure 3-22 Comparison between SCM and MCM

Each SCM is a Ceramic Column Grid Array (CCGA) package where the chip carrier is raised slightly from its board mounting by small metal solder columns that provide the required connections and improved thermal resilience characteristics. As well as the SCM, each processor card also contains the L3 cache and the memory DIMMs, as shown in Figure 3-23. The processor card is mounted in a rugged metal enclosure (book) that protects and secures the card (both in and out of the server), and helps manage airflow used for cooling.




Figure 3-23 POWER4+ processor card layout

Memory access is through the on-chip Level 2 cache (L2) and Level 3 (L3) cache directory controller to the off-chip L3 cache and finally through the memory controller and synchronous memory interface (SMI) to the memory DIMMs, as represented in Figure 3-24 on page 83.




Figure 3-24 Conceptual diagram of POWER4+ processor and memory subsystem

The POWER4+ storage subsystem consists of three levels of cache and the memory subsystem. The first two levels of cache are onboard the POWER4+ chip. The first level is 64 KB of Instruction (I) and 32 KB of Data (D) cache per processor core. The second level is 1.5 MB of L2 cache on the POWER4+ chip. All caches have either full ECC or parity protection on the data arrays, and the L1 cache has the ability to refetch data from the L2 cache in the event of soft errors detected by parity checking.

A 2-way configuration using two 1-way processor cards offers better performance than a configuration using a single 2-way processor card because the maximum capacity of memory is doubled from 16 GB to 32 GB and the Level 2 (L2) and Level 3 (L3) cache on each card is dedicated to a single processor. A 2-way configuration using a 2-way processor card shares the L2 and L3 caches. A comparative table of 2-way configurations is provided in Table 3-10.

Table 3-10 2-way configurations comparative table

2-way configuration	Rperf	Memory GB	L2 cache MB	L3 cache MB
One 2-way 1.2 GHz POWER4+ processor card	3.73	16	1.5 shared	8 shared
L3 Two 1-way 1.2 GHz POWER4+ processor cards	4.18	32	1.5 per processor	8 per processor

2-way configuration	Rperf	Memory GB	L2 cache MB	L3 cache MB
One 2-way 1.45 GHz POWER4+ processor card	4.41	16	1.5 shared	8 shared
Two 1-way 1.45 GHz POWER4+ processor cards	4.94	32	1.5 per processor	8 per processor

Many applications, such as high-performance computing (HPC) and commercial and professional workstation environments, may take advantage of the larger memory capacity and increased L2 and L3 cache per processor available with the two 1-way processor configurations to achieve improved performance. The compute intensive or I/O intensive operations that are typical of these kinds of applications often benefit from large memory capacity and more L2/L3 cache per processor.

The Level 3 (L3) cache consists of two components: The L3 cache controller/directory and the L3 data array. The L3 cache controller/directory is on the POWER4+ chip, and the L3 data array, which consists of 8 MB (POWER4+) of embedded DRAM (eDRAM), is located on a separate module mounted on the processor card.

The conceptual diagram of the memory subsystem of the Model 6C4 and Model 6E4 is shown in Figure 3-24 on page 83. As shown, there are four 8-byte data paths from the memory controller to the memory with an aggregated bandwidth of 6.4 GBps on each processor card. Each processor card can hold up to eight double data rate (DDR) synchronous DRAM (SDRAM) DIMM memory cards, which must be populated in quads.

DDR memory can theoretically double memory throughput at a given clock speed by providing output on both the rising- and falling-edges of the clock signal (rather than just on the rising edge).

The system bus from the processors to the memory subsystem is 2 x 16 bytes at one-third CPU clock speed (or 483.3 MHz on 1.45 GHz POWER4+ systems) for an aggregate data rate of 15.47 GBps. For a system with two processor cards, the fabric bus connects between POWER4+ chips on each card. All traffic to and from the I/O subsystem is through the GX bus on the first processor card. All system, I/O, and PCI-X buses support parity error detection.

The GX controller (embedded in the POWER4+ chip) is responsible for controlling the flow of data through the GX bus. The GX bus is a high-frequency, single-ended, unidirectional, point-to-point bus. Both data and address

information are multiplexed onto the bus, and for each path there is an identical bus for the return path. The GX bus has dual 4-byte paths.

The first processor card connects to the GX bus through its GX controller. The GX bus is connected to a Remote I/O (RIO) bus on the system board through a RIO bridge chip. Each RIO or RIO-2 bus provides 1 byte at 500 MHz or 1 byte at 1 GHz in each direction, or 1 GBps or 2 GBps bidirectionally.

The second processor card has access to the GX bus using a module-to-module fabric bus, which connects it to the first processor card. The fabric bus is similar in nature to the GX bus, but has dual 8-byte paths.

Characteristics of the GX bus and fabric bus are provided in Table 3-11.

Table 3-11 GX bus and fabric bus characteristics

Processor type	GX bus frequency	GX bus data rate	Fabric bus frequency	Fabric bus data rate
1.2 GHz POWER4+	400 MHz	3.2 GBps	600 MHz	9.6 GBps
1.45 GHz POWER4+	483.3 MHz	3.87 GBps	725 MHz	11.6 GBps

The remote I/O bridge contains a series of 1-byte buses grouped in pairs called ports. One of the buses in the pair is for inbound data transfers (500 MBps or 1 GBps), while the other bus is for outbound data transfers (500 MBps or 1 GBps). There is a total of four ports. Two ports are for the internal I/O and two ports are used to attach to any I/O drawers in a loop using RIO cables. The two internal ports connect to a PCI-X Host Bridge (PHB). The PHB chip acts as a bridge between the RIO bus and two PCI-X to PCI-X bridges, which fan out to integrated I/O controllers and slots.

Integrated devices include two Ethernet controllers on the system board and a dual Ultra3 SCSI controller located on the PCI-slot riser card. These integrated devices and I/O slots are further described in the following section.

PCI-X is the latest version of PCI bus technology, using a higher clock speed (133 MHz) to deliver a bandwidth of up to 1 GBps. The PCI-X slots in the Model 6C4 and 6E4 systems support hot-plug and Extended Error Handling (EEH). EEH-capable adapters respond to a specially generated data packet from a PCI slot with a problem. This packet is analyzed by the system firmware, which then allows the device driver to reset the adapter or slot, isolating the error and reducing the need for a system reboot.

The PCI-X slots meet 64-bit 133 MHz requirements, but two slots are affected in performance by an PCI-X to PCI-X bridge that is connected to an ISA bridge, lowering the PCI-X to PCI-X bridge and its connected PCI-X host bridge

specification into 64-bit 66 MHz. PCI adapters that demand high performance should be installed to slots driven by an PCI-X to PCI-X bridge that is not connected to ISA bridge.

The POWER4+ Model 6C4 has six internal PCI-X slots and four disk bays in one 4-pack, which is sufficient for basic operation. If more PCI-X slots and disks are needed, especially well-suited for LPAR mode, up to two 7311-D20 I/O drawers can be added to the Model 6C4. For a Model 6E4, the attachment of the 7311-D20 I/O drawer is not supported.

3.3.4 RAS features

The following are the RAS features of the p630 Model 6C4 and Model 6E4:

- ▶ Fault avoidance through highly reliable component selection and component minimization.
- ▶ Improved reliability through processor operation at a lower voltage enabled by the use of copper chip circuitry and Silicon-on-Insulator technology.
- ▶ Fault tolerance with redundancy, dual line cords, and concurrent maintenance for power and cooling (using optional redundant hot-swap power supplies and fans).
- ▶ Automatic First Failure Data Capture (FFDC) and diagnostic fault isolation capabilities.
- ▶ Concurrent run-time diagnostics based on First Failure Data Capture.
- ▶ Predictive failure analysis on processors, caches, memory, and disk drives.
- ▶ Dynamic error recovery.
- ▶ Error Checking and Correction (ECC) or equivalent protection (such as refetch) on main memory arrays storage, all cache levels (1, 2, and 3), and internal processor arrays.
- ▶ Dynamic processor deallocation based on run-time errors.
- ▶ Persistent processor deallocation (boot-time deallocation based on run-time errors).
- ▶ Persistent deallocation extended to memory.
- ▶ Chipkill correction in memory.
- ▶ Memory scrubbing and redundant bit-steering for self-healing.
- ▶ Industry-leading PCI bus parity error recovery as first introduced on the p690 systems.
- ▶ Hot-plug functionality of the PCI bus I/O subsystem.
- ▶ PCI bus and slot deallocation.

- ▶ Disk drive fault tracking.
- ▶ Avoiding checkstops with process error containment.
- ▶ Environmental monitoring (temperature and power supply).
- ▶ Auto-reboot.
- ▶ Disk mirroring (RAID1) and disk controller duplexing capability are provided by the AIX operating system.
- ▶ Service Agent.
- ▶ NEBS/ETSI compliant (rack mounted server).
- ▶ HACMP.

3.3.5 Minimum and standard features

Table 3-12 summarizes the minimum and standard feature for p630 Model 6C4 and 6E4.

Table 3-12 p630 Model 6C4 and 6E4 minimum and standard features

Features (min./standard)	Description
Microprocessor type	64-bit POWER4, and POWER4+
Min. number of processors	<ul style="list-style-type: none"> ▶ 1-way POWER4+ 1.2 GHz, 8 MB L3 Cache (FC 5133) ▶ 1-way POWER4+ 1.45 GHz, 8 MB L3 Cache (FC 5126)
Memory (minimum)	1 GB
Data/Instruction (L1) cache	64 KB/32 KB
Level 2 (L2) cache	1.5 MB of L2 cache per processor on the 1.2 and 1.45 GHz processor
Internal disk bays	Four available hot-swap disk drive bays
Internal disk drive	Each bay can contain one of the following disks: <ul style="list-style-type: none"> ▶ 36.4 GB Ultra3 10K RPM (FC 3158) ▶ 36.4 GB Ultra320 15K RPM (FC 3277) ▶ 73.4 GB Ultra3 10K RPM (FC 3159) ▶ 73.4 GB Ultra320 15K RPM (FC 3278) ▶ 146.8 GB Ultra320 10K RPM (FC 3275)

Features (min./standard)	Description
Media bays	Two media bays: <ul style="list-style-type: none"> ▶ Bay one for one CD-ROM or DVD-RAM ▶ Bay two for one DVD-RAM, diskette drive, or tape drive (optional)
PCI slots	<ul style="list-style-type: none"> ▶ Six PCI-X slots with 1.2 or 1.45 GHz processor systems ▶ Hot-swap, 64-bit, 133 MHz, and 3.3V
Standard Ports	Three serial ports, two Hardware Management Console ports, one parallel, one keyboard, and one mouse ports
Integrated SCSI Adapters	Two integrated Ultra3 SCSI controllers (one internal and one external)
Integrated LAN adapter	10/100 Mbps Ethernet controller with two ports
System support (service) processor	Yes
Operating system version	<ul style="list-style-type: none"> ▶ AIX 5L Version 5.1 with 5100-02 maintenance package, or later, operating system ▶ Linux 2.4 available from one or more IBM Linux Distribution Partners
Logical Partitioning (LPAR)	Offers greater flexibility in using available capacity and dynamically matching resources to changing business needs.
Warranty	One year

3.3.6 System expansion

Table 3-13 summarizes the possible maximum processor, memory, and storage expansion features for p630 Model 6C4 and 6E4.

Table 3-13 p630 Model 6C4 and 6E4 system expansion features

System expansion features	Description
SMP configuration	<ul style="list-style-type: none"> ▶ 2- or 4-way POWER4+ 1.2 GHz, 8/16 MB L3 Cache. ▶ 2- or 4-way POWER4+ 1.45 GHz, 8/16 MB L3 Cache.

System expansion features	Description
Memory	Up to 32 GB using the following available order combinations: <ul style="list-style-type: none"> ▶ FC 4451, 1024 MB (4 x 256 MB) DIMMs, 208-pin 8 ns DDR SDRAM. ▶ FC 4452, 2048 MB (4 x 512 MB) DIMMs, 208-pin 8 ns DDR SDRAM. ▶ FC 4453, 4096 MB (4 x 1024 MB) DIMMs, 208-pin 8 ns stacked DDR SDRAM. ▶ FC 4454, 8192 MB (4 x 2048 MB) DIMMs, 208-pin 8 ns stacked DDR SDRAM.
Disk drive and disk bay expansion	Up to 587.2 GB (36.4 GB and 73.4 GB drives available). Up to 12 hot-swappable using up to two 7311-D20 I/O drawers (6 bays each).
External Storage	<ul style="list-style-type: none"> ▶ IBM 2104 Expandable Storage Plus (Ultra3 SCSI) ▶ IBM 7133 Serial Disk System (SSA)

3.3.7 Features

Refer to Table 3-31 on page 150 for information about the features for p630 Models 6C4 and 6E4.

3.3.8 Configuration notes

The following section covers the configuration notes for the p630 Model 6C4 and 6E4.

Figure 3-25 on page 90 shows an inside view and external view of the Models 6C4 and 6E4 in order for you to be familiar with the locations of all the various parts and devices.




Figure 3-25 Views of Models 6C4 and 6E4 (POWER4+ system with 6 PCI-X slots)

Processors

The following are the configuration notes regarding the processors:

- ▶ A minimum of one processor card is required.
- ▶ A maximum of four processors on two cards is allowed.
- ▶ The p630 supports a 2-way configuration using a single 2-way processor card (FC 5132, FC 5134, or FC 5127) or two 1-way processor cards (FC 5131, FC 5133, or FC 5126).
- ▶ A 1-way processor card (FC 5131, FC 5126, FC 5133, FC 8106, or FC 8108) cannot be installed with a 2-way processor card (FC 5132, FC 5127, FC 5134, FC 8107, or FC 8109).
- ▶ Processors with different speeds (1.0, 1.2, or 1.45 GHz) cannot be mixed within the system.
- ▶ FC 5126 and FC 5127 require FC 6557, Redundant Cooling

- ▶ Each processor contains a single CPU Single Chip Module (SCM) and that SCM's local memory. This includes:
 - 32 MB of L3 cache on FC 5131 and FC 5132.
 - 8 MB of L3 cache on FC 5126, FC 5127, FC 5133, FC 5134, FC 8106, FC 8107, FC 8108, and FC 8109.
- ▶ Each processor contains 32 KB of data cache and 64 KB of instruction cache.
- ▶ Each 1-way 1.0 GHz processor SCM contains 1.44 MB of L2 cache.
- ▶ Each 1-way 1.2 and 1.45 GHz processor SCM contains 1.5 MB of L2 cache.
- ▶ A 2-way 1.0 GHz processor SCM shares the 1.44 MB of L2 cache.
- ▶ A 2-way 1.2 and 1.45 GHz processor SCM shares the 1.5 MB of L2 cache.
- ▶ 1.2 GHz POWER4+ processors (FC 5133 or FC 5134) require a 6-slot riser card (FC 9556 initial order or FC 6556 MES) and a RIO enabled planar (FC 6575 or FC 9581).
- ▶ 1.45 GHz POWER4+ processors (FC 5126 or FC 5127) require a 6-slot riser card (FC 9556 initial order or FC 6556 MES), a RIO enabled planar (FC 6575 or FC 9581) and a Redundant Cooling Option (FC 6557).
- ▶ Processor upgrades:
 - Processor upgrades are obtained through feature conversions.
 - A 1-way to 2-way configuration upgrade requires a installation of a second 1-way processor card or one feature conversion from 1-way to 2-way processors card.
 - A 2-way configuration using two 1-way processor cards to a 4-way configuration upgrade requires two feature conversions from 1-way to 2-way processors card.

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 32 GB of system memory in a 4-way system. A maximum of 16 GB of system memory in a 1- or 2-way system.
- ▶ A minimum of 1 GB of memory is required in a 1- or 2-way system. A minimum of 2 GB of memory is recommended but not required in a 4-way system.
- ▶ Memory DIMMs must be ordered and installed in quads of equal value (for example 4 x 256 MB DIMMs).
- ▶ Memory plugs into the processor card. Eight memory DIMM slots per processor card. Maximum of two memory features per processor card.

- ▶ Memory should be balanced across the two processor cards (4-way) for best performance.
- ▶ Eight memory DIMM slots are located on the right-hand side of the processor card. Slots 1-4 are on the top of the card and slots 5-8 are on the bottom of the card.
- ▶ A 1-or 2-way system can have a maximum memory of 16 GB. Table 3-14 provides examples of memory DIMM placement.

Table 3-14 Examples of memory DIMM placement

Memory size	DIMM recommended	Feature code	Slot positions
1 GB	4 X 256 MB DIMMs	4451	Slots 1, 3, 6, and 8
2 GB	4 X 512 MB DIMMs	4452	Slots 1, 3, 6, and 8
4 GB	4 X 1024 MB DIMMs	4453	Slots 1, 3, 6, and 8
8 GB	8 X 2048 MB DIMMs	4453	Slots 1, 3, 6, and 8
16 GB	8 X 2048 MB DIMMs	2 X 4454	Slots 1, 3, 4, 5, 6, 7, and 8

A 4-way system can have a maximum memory of 32 GB. Table 3-15 suggests the DIMM positioning for maximum performance.

Table 3-15 Example of memory DIMM placement for maximum performance

Memory size	DIMM recommended	Feature code	Slot positions
1 GB	Not recommended		
2 GB	8 X 256 MB DIMMs	2 X 4451	Slots 1, 3, 6, and 8 in Book 1, and 2
4 GB	8 X 512 MB DIMMs	2 X 4452	Slots 1, 3, 6, and 8 in Book 1, and 2
8 GB	8 X 1024 MB DIMMs	2 X 4453	Slots 1, 3, 6, and 8 in Book 1, and 2
16 GB	8 X 2048 MB DIMMs	2 X 4454	Slots 1, 3, 6, and 8 in Book 1, and 2
32 GB	16 X 2048 MB DIMMs	2 X 4454	Slots 1, 3, 4, 5, 6, 7, and 8 in Book 1, and 2

Figure 3-26 on page 93 shows the memory DIMM slot locations, as well as the LED location associated with each memory DIMM slot.




Figure 3-26 Memory DIMM slot and LED locations

Power supplies

The following are the configuration notes regarding the power supplies:

- ▶ A minimum of one power supply is required on an order.
- ▶ AC system requires an order of FC 6273
- ▶ Redundant power and redundant cooling require:
 - Redundant AC Power (Second FC 6273)
 - Redundant cooling (FC 6557)

Power cords

The following are the configuration notes regarding the power cords:

- ▶ Refer to Appendix C, “Power cord features” on page 907 for further information.

Racks

The following are the configuration notes regarding the racks:

- ▶ A maximum of nine pSeries 630 Model 6C4 can be mounted in a T00 rack.
- ▶ A maximum of ten pSeries 630 Model 6C4 can be mounted in a T42 rack.
- ▶ Supported devices such as SCSI or SSA disk subsystems or tape subsystems may share the rack with the Model 6C4 system.
- ▶ The IBM 7014-T00 rack supports the pSeries 630 with DC power. In order to mount more than four pSeries 630 systems with DC power in a T00 rack, FC 0986 must be specify so the configuration will receive special request services from the Center for Customized Solutions. DC powered pSeries 630 Model 6C4 systems are supported only in the 7014-T00 rack equipped with DC power distribution panel (FC 6117). The 7014-T00 rack power distribution panel (FC 6117) provides attachment for up to four 3-pin to 3-pin PDP power cables (FC 6285) and up to four 3-pin to 5-pin PDP power cables (FC 6286). This will allow a maximum of up to four 6C4s with redundant power or up to seven 6C4s without redundant power. Up to four 7133-D40 (SSA) or 2104-DU3 (SCSI) drawers, or any combination of SSA and SCSI drawers, can also be installed with the four redundant power 6C4s in the same 7014-T00 rack. DC 7133-D40 SSA or 2104-DU3 SCSI drawers are shipped automatically with the correct power cables. Rack FC 6117 provides the proper capacity circuit breakers for use with the pSeries 630 Model 6C4 system based on the inclusion of rack FC 0206 (Rack Content Specify, 7028-6C4) on the 7014-T00 rack order. FC 6117 is a feature of the 7014-T00 rack.
- ▶ The 7311-D20 I/O drawer is designed to be placed at any location in the rack. For rack stability reasons, it is advisable to start filling an empty rack from the bottom and place I/O drawers above system units.
- ▶ The I/O drawers could be in the same rack as the Model 6C4 server or in an adjacent rack, although it is recommended that the I/O drawers be located in the same rack as the server for service considerations.
- ▶ The I/O drawer is an IBM service representative installable item.

I/O drawer attachment

The following are the configuration notes regarding the I/O drawer attachment:

- ▶ For a Model 6E4, the attachment of the 7311-D20 I/O drawer is not supported.
- ▶ Model 6C4 systems can be attached to up to two 7311-D20 I/O drawers.
- ▶ Each 7311-D20 I/O drawer provides seven PCI slots and up to twelve disk drives. See Chapter 5, "IBM rack-mounted expansion drawers" on page 199 for specifications.

- ▶ The Model 6C4 system allows up to two I/O drawers to be attached to the system CEC drawer. I/O drawers should 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 6C4 drawer via the following cables:
 - RIO cables for data transfer
 - 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 6C4 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.
- ▶ RIO ports are provided by the CEC backplane of the Model 6C4 only.
- ▶ Only one RIO loop is available, supporting up to two I/O drawers. A minimum of two RIO cables are required to attach the first I/O drawer on a 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 in various lengths to attach I/O drawers within a single rack or across multiple racks, if desired.
 - FC 3147 = 3.5 meters
 - FC 3148 = 10 meters
- ▶ The configurator validates system connections when the RIO cables are ordered on the individual 7311-D20 drawers. If RIO cables are ordered on the 7028-6C4 server, system validation cannot be performed.
- ▶ Power control for the I/O drawers is provided via one loop. The number of power control (SPCN) cables required is equal to one plus the number of I/O drawers attached to the system. A minimum of two power control cables (FC 6001 = 2 meters, FC 6006 = 3 meters, FC 6007 = 15 meters, or FC 6008 = 6 meters) are required for attachment of the first drawer. Each additional drawer requires one additional power control cable to complete the loop attachment.
- ▶ Each RIO or RIO-2 adapter supports one RIO loop. Each RIO loop supports up to two 7311-D20 I/O drawers or up to four 7311-D10 I/O drawers.
- ▶ To achieve the higher Remote I/O performance of RIO-2, the system must have a RIO-2 Enabled System Planar (FC 9581) and 7311-D20 must have RIO-2 riser adapter (FC 6417). RIO and RIO-2 adapter combinations are supported in the same loop; however, the loop will operate at the standard RIO speed.

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ The Models 6C4 and 6E4 contain four hot-swappable disk drive bays. Each bay may contain one of the following features:
 - 36.4 GB Ultra3 10K RPM (FC 3158)
 - 36.4 GB Ultra320 15K RPM (FC 3277)
 - 73.4 GB Ultra3 10K RPM (FC 3159)
 - 73.4 GB Ultra320 15K RPM (FC 3278)
 - 146.8 GB Ultra320 10K RPM (FC 3275)
- ▶ Two media bays are available (one CD-ROM, DVD-RAM, or DVD-ROM is required). The IDE DVD-ROM or SCSI DVD-RAM can read CD-ROM installation media.
 - Media bay 1 can accommodate IDE or SCSI devices.
 - Media bay 2 can only accommodate SCSI devices or the diskette drive.

Figure 3-27 shows the physical device locations of p630 Model 6C4 and 6E4.




Figure 3-27 p630 Model 6C4 and 6E4 bay locations

Table 3-16 describes the operator panel, diskette, disks, and media bay locations, and SCSI IDs, if applicable.

Table 3-16 Disk and media bay locations and SCSI IDs

Item number	Bay location	Description	SCSI ID
1		Operator Panel	
2	Bay D01	IDE CD-ROM (Default) one IDE DVD-ROM (Optional)	
		Tape Drive (Optional)	SCSI ID 0
3	Bay D02	SCSI DVD-RAM (Optional)	SCSI ID 1
		Diskette Drive	
4	Bay D10	Disk bay	SCSI ID 8
5	Bay D11	Disk bay	SCSI ID 9
6	Bay D12	Disk bay	SCSI ID 10
7	Bay D13	Disk bay	SCSI ID 11

The following methods are supported for operating system boot:

- ▶ CD-ROM, DVD-ROM, and DVD-RAM
- ▶ Internal or external tape drives
- ▶ SCSI disk
- ▶ SSA disk
- ▶ SAN boot using FC 6228 with microcode 3.22A1 or later
- ▶ LAN boot

PCI slots and adapters

The following provides configuration notes regarding the PCI slots and adapters:

- ▶ Six hot-plug PCI-X slots, 64-bit, 133 MHz, 3.3 volt on POWER4+ processor models.
- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of one per FC 6230).
- ▶ The Ultra3 SCSI RAID adapter (FC 2498) cannot be placed in slot 4.

- Feature number 8397 must be placed in slot 3. Slot 4 will be unavailable.
- ▶ Refer to Appendix B, “Adapter placement guidelines” on page 785 for additional I/O adapter configuration notes.

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ Graphics adapter, keyboard, and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported is four.
- ▶ FC 2842 and FC 2843 support only 1- and 2-way configurations.

Table 3-17 provides the maximum quantity of graphic adapters supported for p630 Model 6C4 and 6E4 with four and six slots.

Table 3-17 Graphic adapters supported on systems with 4 and 6 slots

Feature code	Description	Four slots	Six slots
2848	GXT135P Graphics Accelerator	4	4
2842	GXT4500P Graphics Accelerator	1	1
2843	GXT6500P Graphics Accelerator	1	1

Hot-plug options

The following are the configuration notes regarding the hot-plug options:

- ▶ It is not necessary to power down the system to install/remove/replace hot-plug options. Refer to the product documentation for additional information.
- ▶ The following options are hot-pluggable:
 - Hot-plug power supplies
 - Hot-plug disk drive units (DASD)
 - Hot-plug fans
 - Hot-plug adapters

3.3.9 Express configurations

The following are the express configurations available for the Model 6C4 and 6E4. Additional, optional features may be added as desired. Except as noted, all

standard configuration rules for Model 6C4, and 6E4 remain unchanged. Table 3-18 provides a summary of the express configurations for the p630.

Table 3-18 p630 express configuration summary

Express config.	OS	CPU	Memory	DASD
120 C/E	AIX 5L	One 1-way 1.2 GHz POWER4+	2048 MB	Two 36.4 GB
121 C/E	Linux	One 1-way 1.2 GHz POWER4+	2048 MB	Two 36.4 GB
220 C/E	AIX 5L	One 2-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
221 C/E	Linux	One 2-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
222 C/E	AIX 5L	Two 1-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
223 C/E	Linux	Two 1-way 1.2 GHz POWER4+	4096 MB	Two 36.4 GB
420 C/E	AIX 5L	Two 2-way 1.2 GHz POWER4+	8192 MB	Two 36.4 GB
421 C/E	Linux	Two 2-way 1.2 GHz POWER4+	8192 MB	Two 36.4 GB
140 C/E	AIX 5L	One 1-way 1.45 GHz POWER4+	2048 MB	Two 36.4 GB
141 C/E	Linux	One 1-way 1.45 GHz POWER4+	2048 MB	Two 36.4 GB
240 C/E	AIX 5L	One 2-way 1.45 GHz POWER4+	4096 MB	Two 36.4 GB
241 C/E	Linux	One 2-way 1.45 GHz POWER4+	4096 MB	Two 36.4 GB
440 C/E	AIX 5L	Two 2-way 1.45 GHz POWER4+	8192 MB	Two 36.4 GB
441 C/E	Linux	Two 2-way 1.45 GHz POWER4+	8192 MB	Two 36.4 GB

120C for Model 6C4, and 120E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 120C or 120E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (223C or 223E) (FC 8036)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8106)
- ▶ One 2048 MB Memory DIMMS (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)

- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC Power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One Power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One Software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 121C for Model 6C4 and 121E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 121C or 121E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (121C or 121E) (FC 8034)
- ▶ One 1-way POWER4+ 1.2 GHz processor (FC 8106)
- ▶ One 2048 MB Memory DIMMS (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC Power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

220C for Model 6C4 and 220E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 220C or 220E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (220C or 220E) (FC 8177)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8107)
- ▶ Two 2048 MB Memory DIMMS (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 221C for Model 6C4 and 221E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 221C or 221E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (221C or 221E) (FC 8177)
- ▶ One 2-way POWER4+ 1.2 GHz processor (FC 8107)
- ▶ Two 2048 MB Memory DIMMS (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)

- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

222C for Model 6C4 and 222E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 222C or 222E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (222C or 222E) (FC 8035)
- ▶ Two 1-way POWER4+ 1.2 GHz processor (FC 8106)
- ▶ Two 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 223C for Model 6C4 and 223E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 223C or 223E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (223C or 223E) (FC 8036)
- ▶ Two 1-way POWER4+ 1.2 GHz processor (FC 8106)
- ▶ Two 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

420C for Model 6C4 and 420E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 420C or 420E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (420C or 420E) (FC 8037)
- ▶ Two 2-way POWER4+ 1.2 GHz processor (FC 8107)
- ▶ Four 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)

- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 421C for Model 6C4 and 421E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 421C or 421E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (421C or 421E) (FC 8038)
- ▶ Two 2-way POWER4+ 1.2 GHz processor (FC 8107)
- ▶ Four 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

140C for Model 6C4 and 140E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 140C or 140E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (140C or 140E) (FC 8110)
- ▶ One 1-way POWER4+ 1.45 GHz processor (FC 8108)
- ▶ One 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 141C for Model 6C4 and 141E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 141C or 141E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (141C or 141E) (FC 8111)
- ▶ One 1-way POWER4+ 1.45 GHz processor (FC 8108)
- ▶ One 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drives (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)

- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

240C for Model 6C4 and 240E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 240C or 240E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (240C or 240E) (FC 8112)
- ▶ One 2-way POWER4+ 1.45 GHz processor (FC 8109)
- ▶ Two 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other

than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 241C for Model 6C4 and 241E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 241C or 241E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (241C or 241E) (FC 8113)
- ▶ One 2-way POWER4 1.45 GHz processor (FC 8109)
- ▶ Two 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

440C for Model 6C4 and 440E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 440C or 440E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630 (440C or 440E) (FC 8114)
- ▶ Two 2-way POWER4+ 1.45 GHz processor (FC 8109)
- ▶ Four 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)

- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)
- ▶ One software preinstall - (FC 5005)
- ▶ One AIX 5L Version 5.1 or Version 5.2 or later license

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other than processors and memory. Minimum configuration may change depending on features ordered.

Linux Ready 441C for Model 6C4 and 441E for Model 6E4

One 7028-6C4 or 6E4 is part of Express Configuration 441C or 441E, including a minimum of:

- ▶ One feature number indicator for Express Configuration pSeries 630X (441C or 441E) (FC 8115)
- ▶ Two 2-way POWER4+ 1.45 GHz processor (FC 8109)
- ▶ Four 2048 MB memory DIMMs (FC 8080)
- ▶ Two 36.4 GB hot-swap disk drive (FC 3158)
- ▶ One CD-ROM (FC 2633)
- ▶ One Ultra3 SCSI Backplane for hot-swap disk (FC 6568)
- ▶ One SCSI backplane cable (FC 4254)
- ▶ One 6-slot riser indicator (FC 9556)
- ▶ One AC power specify (FC 9172)
- ▶ One AC power supply (FC 6273)
- ▶ One redundant cooling (FC 6557)
- ▶ One language group specify - (FC 9xxx)
- ▶ One power cord specify - (FC 99xx (6C4 only)/FC 98xx)

FC 8xxx are mandatory. Additional, optional 7028-6C4/6E4 features may be added as desired. Upgrades may be made to the minimum configurations, other

than processors and memory. Minimum configuration may change depending on features ordered.

3.4 9114-275 IBM IntelliStation POWER 275

The following is a brief description of the Model 275, as shown in Figure 3-28.



Figure 3-28 IntelliStation POWER 275

3.4.1 General description

The IntelliStation POWER 275 is the first IBM workstation based on the POWER4+, 64-bit, SMP microprocessor, configurable as a 1- or 2-way system running at 1.0 GHz or 1.45 GHz.

The following sections will describe the important technical differences of the pSeries 615.

3.4.2 Configuration notes

The following are the known configuration notes at the time of writing.

Processors

The IntelliStation POWER 275 is orderable as a 1- or 2-way POWER4+ microprocessor running at 1.0 GHz or 1.45 GHz and supports 8 MB of L3 cache.

Graphics adapters

The GXT4500P Graphics Accelerator (FC 2842) is a 64-bit entry 3D PCI Graphics adapter. The graphics subsystems provide excellent functionality and performance for demanding graphics applications. This adapter has the following base features:

- ▶ 128 MB Unified Frame Buffer
 - 24-bit Double Buffered up to 2048 x 1536
 - 24-bit Double Buffered Stereo up to 1280 x 1024
 - 24-bit Z-Buffer
 - 4/8-bit Overlay
 - 8-bit Double Buffered Alpha
 - 8-bit Stencil
 - Eight windows ID bits
 - Four Clipping Planes
- ▶ Texture Mapping
 - Up to 110 MB Texture Memory (1280 x 1024)
 - Dual Texture
 - 3D texture
- ▶ Four Hardware Color Maps

The GXT6500P Graphics Accelerator (FC 2843) is a 64-bit mid-range 3D PCI Graphics adapter. The graphics subsystems provide excellent functionality and performance for demanding graphics applications. This adapter has the following base features:

- ▶ 128 MB Unified Frame Buffer
 - 24-bit Double Buffered up to 2048 x 1536
 - 24-bit Double Buffered Stereo up to 1280 x 1024
 - 24-bit Z-Buffer

- 4/8-bit Overlay
- 8-bit Double Buffered Alpha
- 8-bit Stencil
- Eight windows ID bits
- Four Clipping Planes
- ▶ Texture Mapping
 - Up to 110 MB Texture Memory (1280 x 1024)
 - Dual Texture
 - 3D texture
- ▶ Four Hardware Color Maps
- ▶ Full OpenGL, graPHIGS Geometry Accelerator

Audio support

The Audio PCI Adapter (FC 8244) for workstations is a high-performance workstation class audio PCI adapter that is optimized for support in the AIX professional workstation environment and their associated applications. The Audio PCI Adapter for workstations provides external jacks for headphones, speaker output, line input, and microphone input, as well as an internal connector for CD or DVD drive audio input.

3.4.3 Express configurations

The following are the available express configurations available for the 9114-275.

Express configuration 140W (FC 8181)

The 140W express configuration for the IntelliStation Model 275 consists of the following key features:

- ▶ One 1-way 1.45 GHz POWER4+ processor (FC 8182)
- ▶ One 2048 MB (4x512 MB) DIMM (FC 8157)
- ▶ One 36.4 GB 10,000 RPM Ultra320 SCSI disk drive (FC 3273)

Express configuration 240W (FC 8183)

The 240W express configuration for the IntelliStation Model 275 consists of the following key features:

- ▶ One 2-way 1.45 GHz POWER4+ processor (FC 8184)
- ▶ One 4096 MB (4x1024 MB) DIMM (FC 8158)
- ▶ Two 36.4 GB 10,000 RPM Ultra320 SCSI disk drive (FC 3273)

Express configuration 100W (FC 8186)

The 100W express configuration for the IntelliStation Model 275 consists of the following key features:

- ▶ One 1-way 1.0 GHz POWER4+ processor (FC 8185)
- ▶ One 2048 MB (4x512 MB) DIMM (FC 8157)
- ▶ One 36.4 GB 10,000 RPM Ultra320 SCSI disk drive (FC 3273)

3.5 7043 Model 150 IBM RS/6000 43P

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

Check the latest sales information for system availability.

The Model 150 is a uni-processor system that provides enhanced performance over its predecessor, the Model 140, by utilizing a 250 or 375 MHz PowerPC 604e processor and an enhanced memory controller. With this memory controller, the Model 150 uses SDRAM memory and an 83 MHz memory bus speed. The system memory can also be expanded up to 1 GB.

With Ethernet and Ultra SCSI controllers integrated on the planar, the Model 150 also contains five PCI slots and five bays for expansion and growth capability. This makes the Model 150 an attractive investment as either a technical workstation or an entry workgroup server.

The Model 150 supports a variety of 2D and 3D graphics adapters (including a new advanced 3D graphics adapter), offering excellent graphics price and performance. In addition, a robust set of disk drive and communications features is available.

The AIX 5L Version 5.1, or later, operating system is included with each Model 150 and can be preinstalled.

Figure 3-29 on page 113 shows the RS/6000 Model 150.




Figure 3-29 Model 7043-150 with peripherals

3.5.1 Product positioning

The Model 150 is an entry-level to mid-range technical workstation or entry-level workgroup server.

The Model 150 is the high end of the 43P family, since its additional function and advanced graphics capability offer a significant price/performance improvement over the 332 MHz Model 140.

3.5.2 Highlights

The Model 150 provides an excellent price/performance solution in the entry RS/6000 graphics workstation and workgroup families. The Model 150 offers:

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

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

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

3.5.3 Technical overview

The RS/6000 Model 150 is a uni-processor system that utilizes the PowerPC 604e processor running at 250 MHz or 375 MHz. It supports an integrated 1 MB parity-checked synchronous L2 cache. The Model 150 also features an enhanced memory controller that uses ECC SDRAM memory.

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

Figure 3-30 on page 115 shows the system diagram for the RS/6000 Model 150.




Figure 3-30 RS/6000 Model 150 diagram

In the Model 150, the 60X bus runs from the PowerPC 604e processor and connects to an enhanced memory controller chip, which has a processor system interface of a 64-bit data bus and a 32-bit address bus. This highly integrated chip acts as both a PCI bridge and a memory controller. It contains the system bus arbitration, provides support for full memory coherency, and pipelining of processor accesses.

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

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

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

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

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

3.5.4 RAS features

The following are the RAS features of the RS/6000 Model 150:

- ▶ ECC (Error-Checking and Correcting)
- ▶ Fan speed monitoring
- ▶ AC power loss sensing
- ▶ Run-Time Error Capture capability
- ▶ Online (concurrent) Diagnostics with Error Log Analysis and Service Aids
- ▶ Auto-restart (reboot) option
- ▶ System Management Service (SMS)
- ▶ Service Agent

3.5.5 Minimum and standard features

This section describes the system minimum and standard features for RS/6000 Model 150.

Table 3-19 summarizes the minimum and standard features.

Table 3-19 RS/6000 Model 150 minimum and standard features

Features (min./standard)	Description
Microprocessor	250 MHz or 375 MHz PowerPC 604e
Level 1 (L1) cache	32 KB data/32 KB instruction

Features (min./standard)	Description
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	36.4 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-in. diskette drive
Service Processor	No
AIX operating system version	AIX 5L Version 5.1 operating system or later
Warranty	24 x 7, on-site for one year (limited), at no additional cost

3.5.6 System expansion

This section describes the available features for RS/6000 Model 150 expansion.

Table 3-20 shows the possible storage and memory expansion configurations.

Table 3-20 Model 150 system expansion

Model 150 system expansion	
RAM	Up to 1 GB of SDRAM
Internal disk storage	Up to 91 GB
External disk storage	For additional information, see Appendix E, "Supported peripherals by device matrix" on page 931

3.5.7 Features

Refer to Table 3-31 on page 150 for information about the features for the 7044 Model 150.

3.5.8 Configuration notes

The following section covers the configuration notes for the 7044 Model 150.

Processors

The following are the configuration notes regarding the processors:

- ▶ The following processor cards are supported on this system:
 - 1-way 250 MHz PowerPC 604e with 1 MB of L2 cache
 - 1-way 375 MHz PowerPC 604e with 1 MB of L2 cache

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 1 GB of system memory.
- ▶ A minimum of 128 MB of memory is required
- ▶ One to four memory DIMMs (64 MB, 128 MB, or 256 MB) can be installed (the DIMMs do not have to be installed in pairs).
- ▶ The DIMMs have to be installed in memory slots J11, J28, J19, and J139.

Figure 3-31 on page 119 shows the RS/6000 Model 150 system board as a reference for the memory card locations.




Figure 3-31 System board locations of RS/6000 Model 150

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ Bay 3 can accommodate a second optical storage device capable of reading a CD-ROM disc, tape drive, or other media device.
- ▶ A 4.5 GB (FC 2900), 9.1 GB (FC 2908 or FC 3027), or 18.2 GB (FC 2909 or FC 3102) hard disk can be installed in Bay 3 with a mounting kit (FC 6509).
- ▶ All DASD and media devices are driven by the internal/integrated SCSI port only.

- ▶ Refer to Figure 3-37 and Table 3-24 for additional configurations for storage devices.

Figure 3-32 shows the front view of the RS/6000 Model 150 as reference for the disks and media bay locations.




Figure 3-32 RS/6000 Model 150 front view

Table 3-21 describes the disks and media bay locations and the SCSI-IDs.

Table 3-21 Disk and media bay locations and SCSI-IDs

Bay location	Description	SCSI ID
Bay 1	Diskette drive	Non-SCSI
Bay 2	CD-ROM	SCSI ID 3
	DVD-RAM	Non-SCSI
Bay 3	Optional storage device	SCSI ID 2, or 6
Bay 4	Disk drive	SCSI ID 4
Bay 5	Disk drive	SCSI ID 5

Note: The SCSI bus IDs listed are the recommended values. The SCSI IDs shown for media devices indicate how the devices are set when shipped from the factory. Field installations might not comply with these recommendations.

Boot is supported through:

- ▶ CD-ROM/DVD-ROM/DVD-RAM
- ▶ Internal or external tape drives
- ▶ SCSI disk

- ▶ SSA disk
- ▶ LAN boot

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ Slots 1 and 5 are short slots.
- ▶ Slots 2, 3, and 4 are long slots.
- ▶ Slots 2 and 3 are connected to the primary PCI bus on the planar.
- ▶ Slots 1, 4, and 5 are on the secondary PCI bus provided by the bridge chip on the riser card.
- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of 1 per FC 6230).
- ▶ Refer to Appendix B “Adapter Placement Guidelines” for additional I/O adapter configuration notes.

Figure 3-33 on page 122 shows the PCI slot locations on the riser card of the RS/6000 Model 150.

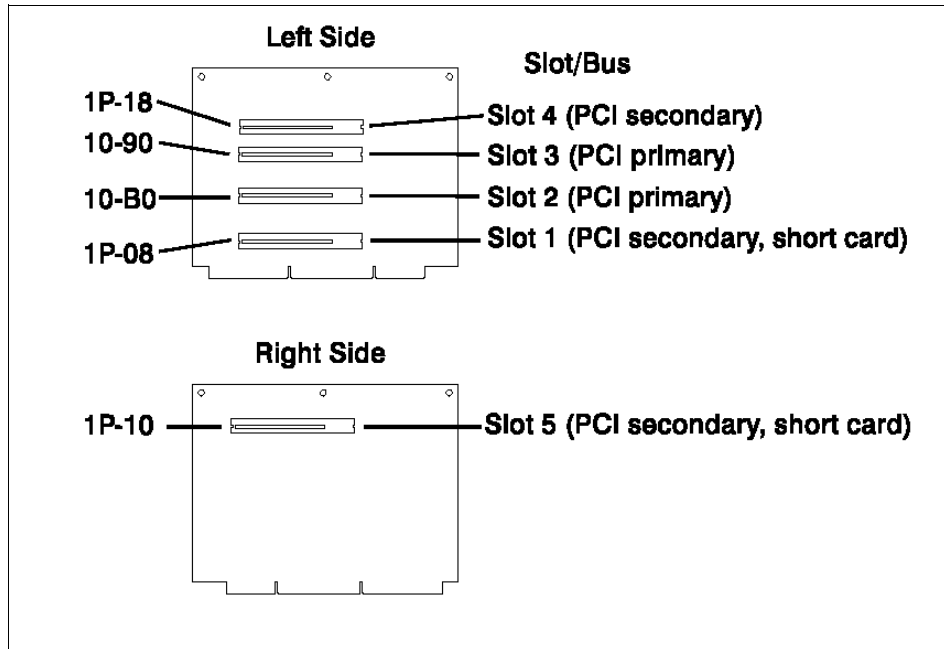


Figure 3-33 PCI riser card for RS/6000 Model 150

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ Graphics adapters have a higher priority than I/O adapters, except for FC 4951, which has the highest priority.
- ▶ The maximum number of graphics adapters is a total of four, in any combination.
- ▶ When FC 4951 is installed in either slot 2 or 3, FC 2823, FC 2841, FC 2851, and FC 2852, if installed, must be installed in slots 1, 4, or 5.
- ▶ When FC 6205 is installed, place FC 2823 in slots 1, 4, or 5.

3.6 7044 Model 170 IBM RS/6000 44P

The IBM RS/6000 44P 7044 Model 170 is an entry-level desktop RS/6000 workstation or workgroup server offered at an affordable price. The Model 170 provides a continuation of the successful line of entry workstations, offering enhanced performance over the 43P 7043 Model 150. The Model 170 is an ideal solution for those who require high-performance 2D or 3D CAD/CAM workstations in a compact but flexible package.

The following features are included in the Model 170:

- ▶ Power3-II processor for performance upgrade options
- ▶ An enhanced memory controller that uses SDRAM memory and a 95 MHz memory bus speed on the 333 MHz system, and a 100 MHz memory bus speed on the 400 or 450 MHz systems
- ▶ Memory expansion capabilities up to 2 GB
- ▶ IBM's high-performance 3D accelerators to meet the needs of demanding design and engineering workstation users
- ▶ Integrated Ethernet, Ultra2 SCSI, and a service processor
- ▶ Six PCI slots and six bays for expansion and growth capability
- ▶ Support of a variety of 2D and 3D graphics adapters, offering excellent graphics price and performance
- ▶ A powerful set of available disk drive and communications features

Check with your IBM representative for the availability of this product.

The AIX 5L Version 5.1, or later, operating system is included with each Model 170 and can be preinstalled, if desired.

The following sections outline the standard and optional features of the Model 170, as shown in Figure 3-34 on page 124.



Figure 3-34 Model 7044-170 with peripherals

3.6.1 Product positioning

The Model 170 is an entry-level to mid-range technical workstation for the RS/6000 platform. It is also an entry-level workgroup server. The 7044-170 may be viewed as an extension/follow-on to the popular 7043-140 and 7043-150 systems.

The Model 170's additional function and advanced graphics capability offers a significant price/performance improvement over the current 375 MHz 7043-150.

Graphics workstation customers

The Model 170 provides support for a wide range of graphics adapters, from the 2D POWER GXT135P to the 3D GXT4500P or GXT6500P. All of these graphics adapters make the Model 170 an ideal solution if you need a high-performance 2D or 3D workstation with expansion capabilities within a single package.

With a choice of high-function 3D graphics accelerators, the Model 170 is well suited for demanding 3D applications, including MCAD engineering analysis, petroleum, and seismic interpretation.

Workgroup Server Customers

The Model 170 can be deployed as a workgroup server or stand-alone business system, providing functions that includes Internet, database, application, and print/file serving.

The wide range of applications for e-business, e-commerce, ERP, and supply-chain management makes the Model 170 an excellent choice for many small-to medium-size businesses. It provides an attractive price performance option over the Model 140 or 150 while expanding the capacity of memory to 2 GB and the slots and bays to six each. It also provides a higher performance option for F40 or F50 customers who do not require the expansion capacity or require a small package.

3.6.2 Highlights

The RS/6000 44P Models 170 extends the IBM line of powerful and affordable workstation/workgroup servers with 64-bit, copper POWER3-II processors. The Model 170 is an excellent 2D or 3D workstation or a powerful workgroup server.

- ▶ General features and benefits
 - Excellent price/performance solution
 - POWER3-II microprocessor (choice of three speeds: 333 MHz, 400 MHz, or 450 MHz)
 - 64-bit system architecture
 - Concurrent 32-bit and 64-bit application support
 - Built-in service processor
 - Ultra2 SCSI disk support
 - SSA RAID disk support
 - AIX operating system
- ▶ Workstation benefits
 - Top performing workstation and 3D graphics for CAD/CAM and petroleum visualization applications
 - Twice the memory of the Model 150 to support larger MCAD and CAE models
 - Choice of 3D graphics accelerators
- ▶ Server benefits
 - Powerful, affordable 64-bit workgroup or e-business server
 - Six bays and six slots, and up to 2 GB of memory

- ▶ IBM SystemXtra® support services and financing

3.6.3 Technical overview

This section provides the technical overview of the 7044 Model 170 and the diagram of the system architecture.

The Model 170 configuration includes a single POWER3-II 64-bit processor with a choice of either a 333 MHz processor with 1 MB of Level 2 (L2) cache or a 400 MHz processor with 4 MB of L2 cache. Currently, the 450 MHz processor option is the fastest POWER3-II processor available in the RS/6000 product family. Each processor is equipped with 64 KB of data and 32 KB of instruction Level 1 (L1) cache. A 10/100 Mbps Ethernet controller, one Ultra SCSI controller, one external Ultra2 SCSI controller, and the service processor are integrated on the system planar.

In addition, the Model 170 has 256 MB of ECC synchronous dynamic random access memory (SDRAM), three disk bays, three media bays, six PCI slots, a CD-ROM drive, a 1.44 MB 3.5-in. diskette drive, and an operator panel. The operator panel has a 2 x 16 backlit LCD for system status and diagnostic information. Microphone and headphone jacks are built into the operator panel. The following ports are included: keyboard, mouse, Ethernet (AUI and RJ45), parallel, two serial (9-pin D-shell), tablet (for use with legacy input devices), Ultra2 SCSI VHDC, and stereo/audio.

Figure 3-35 on page 127 shows the 7044 Model 170 diagram.

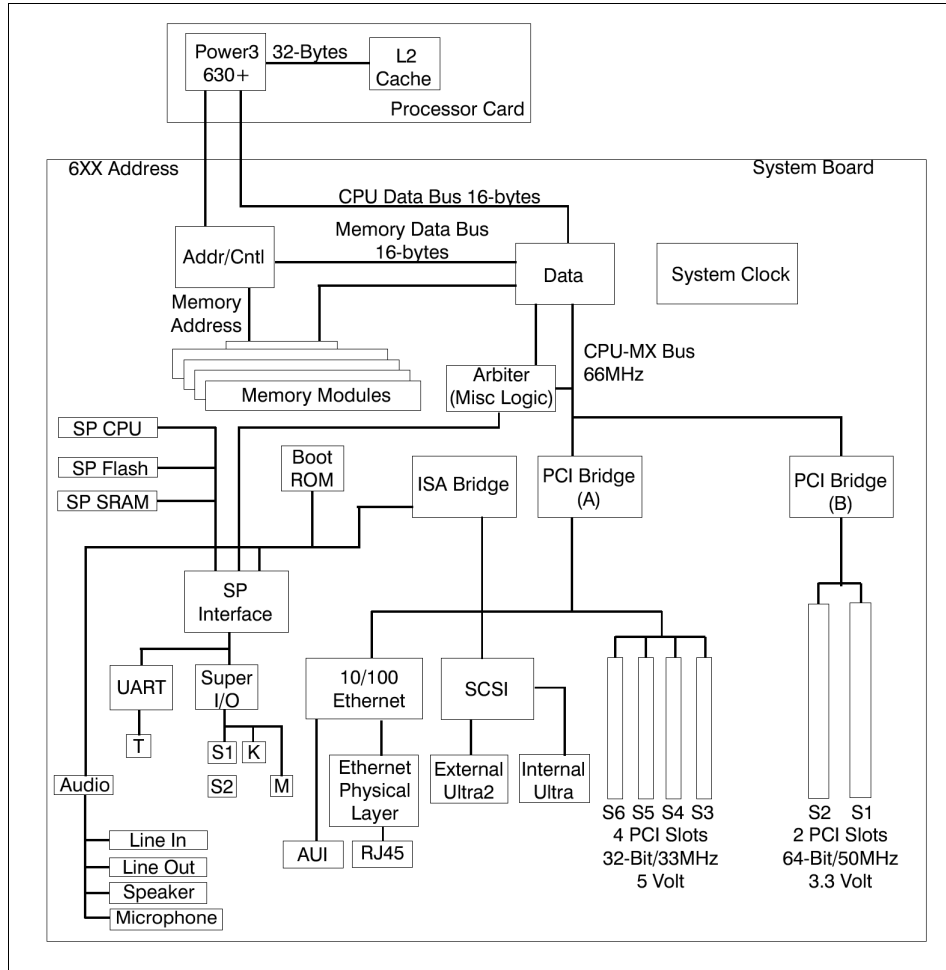


Figure 3-35 System logic flow diagram for 44P Model 170

The 44P Model 170 is the most affordable 64-bit system that IBM has introduced and uses the POWER3-II processor. The POWER3-II processor is designed for high-performance commercial and graphical computing applications. The processor allows for concurrent operation of up to eight instructions per cycle (three fixed-point instructions, two floating point instructions, two load/store instructions, and one branch instruction). Increasing the number of simultaneously executed instructions results in better system performance, especially with applications that are optimized to take advantage of this capability.

The Model 170 uses a 64 KB data and 32 KB instruction 128-way set associative L1 cache. The size of both data and instruction cache reduces the number of cache misses, results in more cache hits, and maximizes performance. Both the data and instruction cache are parity protected.

The 44P Model 170 processor card has either 1 MB (333 MHz), 4 MB (400 MHz) or 8 MB (450 MHz), of L2 cache located on the processor card. L2 cache is used to lower the time spent accessing memory data and increase performance. The L2 cache extends L1 cache benefits by adding more cache to the memory pipeline.

The L2 cache uses a direct mapped cache methodology. There is a dedicated external interface to the L2 cache not shared with the 6XX bus. This allows concurrent access to both the L2 cache and the 6XX bus.

The system supports from 256 MB to 2 GB of ECC SDRAM using a 128-bit wide memory bus. Memory bus speed varies dependent upon processor selection. The enhanced memory controller uses a 95.1 MHz memory bus speed on the 333 MHz system and a 100 MHz memory bus speed on the 400 MHz and 450 MHz systems.

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 (memory bus), and the I/O (I/O bus). The 6XX bus is a 128-bit bus running at a clock speed of either 100 MHz, a 4:1 ratio, when featured with a 400 MHz or 450 MHz processor (which results in a peak data throughput of 1600 MB per second) or 95.1 MHz, a 7:2 ratio, when featured with a 333 MHz processor.

The 6XX bus 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 mode and are pipelined so that new requests may be issued before previous requests are completed.

The Model 170 is compliant with Revision 2.1 of the peripheral component interconnect (PCI) specifications and implements dual PCI bridge chips in a peer configuration. One PCI bridge chip provides a 32-bit interface operating at 33 MHz for four PCI slots and the other PCI bridge chip implements a 64-bit bus operating at 50 MHz for two PCI slots. Slots one and two are 64-bit, 50 MHz, 3.3V slots, and slots three, four, five, and six are 32-bit, 33 MHz, 5.0V slots. All

slots in the Model 170 accept full-sized PCI adapters. The 64-bit slots are physically keyed to accept either universal or 3.3V cards only. 5.0V cards will not seat in the card slots.

3.6.4 RAS features

The following are the RAS features of the RS/6000 Model 170:

- ▶ ECC (Error-Checking and Correcting)
- ▶ Fan speed monitoring
- ▶ AC power loss sensing
- ▶ Run-Time Error Capture capability
- ▶ Online (concurrent) Diagnostics with Error Log Analysis and Service Aids
- ▶ Auto-restart (reboot) option
- ▶ System Management Service (SMS)
- ▶ Service Processor
- ▶ Service Agent

3.6.5 Minimum and standard features

Table 3-22 lists the RS/6000 Model 170 minimum and standard features.

Table 3-22 Model 170 minimum and standard features

Features (min./standard)	Description
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	36.4 GB Ultra SCSI
Disk/media bays	Six (three available)
Expansion slots	Six PCI

Features (min./standard)	Description
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-in. diskette drive
Service Processor	Standard (on planar)
AIX operating system version	AIX Version 5.1, or later, operating system
Warranty	On-site for one year (limited) at no additional cost

3.6.6 System expansion

Table 3-23 shows the possible processor, storage, and memory expansion configurations.

Table 3-23 Model 170 system expansion

Model 170 system expansion	
Microprocessor	400 or 450 MHz POWER3-II
Level 2 (L2) Cache	4 or 8 MB
RAM (memory)	Up to 2 GB
Internal disk storage	Up to 293.6 GB
External disk storage	IBM 2104 Expandable Storage Plus (Ultra2 SCSI), IBM 7133 Serial Disk System (SSA)

3.6.7 Features

Refer to Table 3-31 on page 150 for information about the features for 7044 Model 170.

3.6.8 Configuration notes

The following section covers the configuration notes for the 7044 Model 170.

Processors

The following are the configuration notes regarding the processors:

- ▶ The following processor cards are supported on this system:
 - 1-way 333 MHz POWER3-II with 1 MB of L2 cache
 - 1-way 400 MHz POWER3-II with 4 MB of L2 cache
 - 1-way 450 MHz POWER3-II with 8 MB of L2 cache

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 2 GB of system memory.
- ▶ A minimum of 256 MB of memory is required.
- ▶ One to four memory DIMMs (128 MB, 256 MB, or 512 MB) can be installed (the DIMMs must be installed in pairs).

Figure 3-36 shows the memory card locations.




Figure 3-36 Memory card locations

Note: Memory DIMMs must be installed in pairs and in the correct connector configuration. Connector ID numbers are labeled on the system board. Possible connector pair combinations are J1 and J2, J9 and J10.

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ Bay 2 can accommodate a second optical storage device capable of reading a CD-ROM, tape drive, or other media device.
- ▶ When installing an internal 3.5" disk in Bay 2, use mounting kit FC 6561 for both 10k rpm and 7200 rpm drives. Mounting kit FC 6560 is used for 7200 rpm drives only.
- ▶ Mounting kit (FC 6561) is to be used for installing 10K rpm hard disks in Bay 2.
- ▶ All DASD and media devices are driven by the internal/integrated SCSI port only.
- ▶ Refer to Figure 3-35, and Table 3-28 below for additional configurations for storage devices.

Figure 3-37 on page 133 shows the disks and media bay locations.




Figure 3-37 Disks and media bay locations

Table 3-24 describes the disks and media bay locations and the SCSI IDs.

Table 3-24 Disk and media bay locations and SCSI-IDs

Item number	Bay location	Description	SCSI ID
1	Bay D01	CD-ROM	SCSI ID 1
2	Bay D02	Media device	SCSI ID 2
3	Bay D03	Diskette drive	Non-SCSI
4	Bay D04	Disk drive	SCSI ID 4
5	Bay D05	Disk drive	SCSI ID 5
6	Bay D06	Disk drive	SCSI ID 6

Note: The SCSI bus IDs listed are the recommended values. The SCSI IDs shown for media devices indicate how the devices are set when shipped from the factory. Field installations might not comply with these recommendations. SCSI ID 7 is reserved for diagnostic purposes.

Boot is supported using one of the following methods:

- ▶ CD-ROM/DVD-ROM/DVD-RAM
- ▶ Internal or external tape drives
- ▶ SCSI disk
- ▶ SSA disk
- ▶ SAN boot using a 2 GB Fibre Channel Adapter FC 6239
- ▶ LAN boot

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ Slots 1 and 2 are 64-bit and run at 50 MHz.
- ▶ Slots 3, 4, 5, and 6 are 32-bit and run at 33 MHz.
- ▶ 32-bit slot is half-size.
- ▶ 64-bit slots are full-size
- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of 1 per FC 6230).
- ▶ Refer to Appendix B “Adapter Placement Guidelines” for additional I/O adapter configuration notes.

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ The graphics adapter, keyboard, and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported in the 7044 Model 170 is four.

3.7 9112-265 IBM IntelliStation POWER 265

The IBM 9112 IntelliStation POWER Workstation is a high-performance desktide/desktop workstation that is a follow-on to the highly successful RS/6000 7044 Model 270. It is an affordable 64-bit symmetric multiprocessing system (SMP) system that offers a significant price/performance improvement over the Model 270.

Figure 3-38 shows the IBM IntelliStation POWER 265.



Figure 3-38 IBM IntelliStation 9112 Model 265

The IntelliStation POWER Workstation lowers the cost of a high-end design and analysis applications workstation. The availability of 2D and 3D graphics adapters position the IntelliStation POWER Workstation as a single-seat MCAD design and analysis solution.

The symmetric multiprocessor (SMP) uses one or two state-of-the-art, 64-bit, copper-based, POWER3-II microprocessors running at 450 MHz. Each 450 MHz processor includes 4 MB of Level 2 (L2) cache. The base 512 MB of main memory can be expanded to 8 GB for faster performance and exploitation of 64-bit addressing.

The IntelliStation POWER Workstation contains four bays, expandable to 10. The bays are used for a diskette drive, a CD-ROM or a DVD-RAM, a non-hot-swap disk, and an optional media device. A feature option will allow you to add six front-accessible, hot-swap capable bays that can accommodate up to 440.4 GB of disk storage, for a total of 587.2 GB of internal disk storage.

Light Path Diagnostics are also built into the IntelliStation POWER Workstation LEDs (physically located on key system components) assist in quick diagnosis and resolution of problems, should they arise.

3.7.1 Product positioning

With its performance, SMP capability, and advanced graphics the POWER 265 is well positioned in the mid-range and high-end MCAD workstation price point. Its enhanced floating-point capability and overall computing capability extends opportunities by providing a single seat MCAD design and analysis solution. The IntelliStation POWER Workstation model 265 is at the high end of our workstation offerings, replacing the Model 270 as our design and analysis workstation.

The IntelliStation POWER Workstation model 265's 2-way SMP capability and up to 8 GB of memory make it an excellent solution for demanding technical applications. It replaces the Model 270 as a deskside/desktop technical server.

3.7.2 Highlights

The IntelliStation POWER Workstation Model 265 offers the following:

- ▶ Outstanding value.
- ▶ 3D SMP workstation.
- ▶ Excellent performance for 3D CAD/CAM and petroleum visualization applications.
- ▶ Powerful analysis workstation for complex computational chemistry, computational fluid dynamics, and structural analysis applications.
- ▶ Up to two, state-of-the-art, IBM 64-bit POWER3-II processors running at 450 MHz that exploit unique copper-based technology for high performance and reliable service.
- ▶ Four bays and five PCI slots for feature expandability. Six additional hot-swap DASD bays available for a total of 587.2 GB of internal disk storage.
- ▶ Light Path Diagnostics.

3.7.3 Technical overview and RAS features

Note: For technical information of this product, refer to the pSeries 610 Model 6E1.

3.7.4 Minimum and standard features

Table 3-25 on page 137 summarizes the minimum and standard feature for IntelliStation Model 265.

Table 3-25 IntelliStation Model 265 minimum and standard features

Features (min./standard)	Description
Microprocessor type	64-bit POWER3-II
Min. number of processors	1-way 450 MHz, 8 MB L2 Cache (FC 5323)
Memory (minimum)	512 MB
Memory bus width	Quad 128-bit
Data/Instruction (L1) cache	64 KB/32 KB
Level 2 (L2) cache	4 MB of L2 cache
Internal disk bays	Six available hot-swap disk drive bays and two non hot-swappable disk drives in the media bays
Internal disk drive	36.4 GB Ultra SCSI
Media bays	Four media bays: One diskette drive, One CD-ROM or DVD-ROM, One non hot-swap disk drive (default), One CD-ROM, DVD-ROM, disk drive, or tape drive (optional)
PCI slots	Five PCI
PCI bus width	32- and 64-bit
Standard Ports	Three serial ports, one parallel, one keyboard, and one mouse ports
Integrated SCSI Adapters	Two integrated Ultra3 SCSI controllers (one internal and one external)
Integrated LAN adapter	10/100 Mbps Ethernet controller with two ports
System support (service) processor	Yes
Operating system version	AIX Version 4.3.3 and AIX 5L Versions 5.1 and 5.2
Warranty	One year

3.7.5 System expansion

Table 3-26 summarizes the possible maximum processor, memory, and storage expansion features for p610 Model 6C1 and 6E1.

Table 3-26 IntelliStation Model 265 system expansion features

System expansion features	Description
SMP configuration	2-way 450 MHz, 8 MB L2 Cache (FC 5309)
Memory	Up to 8 GB
Internal disk drive	36.4 GB, 73.4 GB, or 146.8 drives available
External storage	IBM 2104 Expandable Storage Plus (Ultra3 SCSI), IBM 7133 Serial Disk System (SSA)

3.7.6 Features

Refer to Table 3-31 on page 150 for information about the features for IntelliStation Model 265.

3.7.7 Configuration notes

This section reports only the differences between the p610 and the Model 265. For further information, refer to the configuration notes of the p610 Model 6C1 and 6E1.

Processors

- ▶ Processor cards must be installed starting in slot 1 first, then slot 2.

Graphics adapters

- ▶ The keyboard and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported in the IntelliStation POWER Workstation is four.

Table 3-27 on page 139 reports the maximum quantity of graphic adapters supported for the IntelliStation POWER 265.

Table 3-27 Maximum number of graphic adapters supported on systems

Feature code	Description	Quantity max.
2848	GXT135P Graphics Accelerator	4
2842	GXT4500P Graphics Accelerator	4
2843	GXT6500P Graphics Accelerator	2

3.7.8 Express configurations

At the time of writing, there is no Express configuration available for the IntelliStation Model 265.

3.8 7026 Model B80 IBM @server pSeries 640

The IBM @server pSeries 640 (7026-B80) is a rugged system that is ideal for environments such as ISP/ASP, e-commerce, telco/wireless, or scientific and technical computing, where high density, rapid horizontal growth, and leading-edge UNIX performance are critical.

Check with your IBM representative for the availability of this product.

Figure 3-39 shows the pSeries 640 Model B80.



Figure 3-39 pSeries 640 Model B80 entry rack server

3.8.1 Product positioning

With its performance, SMP capability, and high-density packaging, the pSeries 640 is well positioned for transaction-intensive server solutions, such as e-commerce, database, and as an application server. The pSeries 640 is also well suited for operations in telecommunications and wireless networks. The pSeries 640 provides 4-way SMP capability, up to 16 GB of memory, and up to 293.6 GB of internal DASD, which makes it a solution for many demanding commercial applications. It is a cost-effective performance alternative to the RS/6000 7025 Models F50 or F80, or 7026 Models H70 or H80 for commercial customers who do not require the expansion capabilities that these models offer.

The pSeries 640 also offers a variety of high-availability features, such as:

- ▶ An integrated service processor
- ▶ Redundant cooling and redundant power supply (AC or DC)
- ▶ Processor deallocation
- ▶ Hot-swap cooling fans and disk drives
- ▶ A programmable beacon that can be triggered remotely to quickly alert an on-site technician

3.8.2 Highlights

The pSeries 640 Model B80 delivers a cost-efficient growth path to the future through such attributes as:

- ▶ Outstanding value in a powerful, symmetric multiprocessing (SMP), e-business server optimized for high-density environments
- ▶ Rugged, space-saving rack-mount, slide-in chassis packaging designed for extreme levels of availability, maintainability, and ease-of-use with tool-less access
- ▶ Design optimized for ease-of-growth by providing internal upgrade paths and horizontal scaling options to meet rapidly changing capacity demands
- ▶ Up to four 64-bit POWER3-II processors that exploit unique copper-based technology for high-performance reliable service with the following processor options:
 - 1-way or 2-way 375 MHz processors with 4 MB of Level 2 (L2) cache per processor
 - 2-way 375 MHz processor with 8 MB of L2 cache per processor
 - 2-way 450 MHz processor with 8 MB of L2 cache per processor

- ▶ Up to 16 GB of memory and 293.6 GB of internal disk storage for the performance and capacity needed by demanding e-business applications
- ▶ Innovative systems management feature using handheld devices such as the Palm Pilot
- ▶ NEBS Level 3 compliant for the telecommunications carrier grade environment
- ▶ Multiple power options, including redundant 110V/220V AC or redundant -48V DC

3.8.3 Technical overview

This section provides the technical overview of the p640 Model B80 and the diagram of the system architecture.

The model B80 5 EIA units (5U) x 19-in. industry-standard rack drawer is 24 inches deep and contains five bays, one of which is used for a 5.25-in. SCSI CD-ROM, DVD-RAM, or tape drive. The remaining four front-accessible, hot swappable bays can accommodate up to 293.6 GB of internal disk storage; an 36.4 GB Ultra SCSI drive is included in the base configuration. Five PCI slots are provided, along with an integrated 10/100 Mbps Ethernet controller with two ports, internal and external Ultra2 SCSI controllers, and a service processor.

Figure 3-40 on page 142 shows the p640 system architecture.

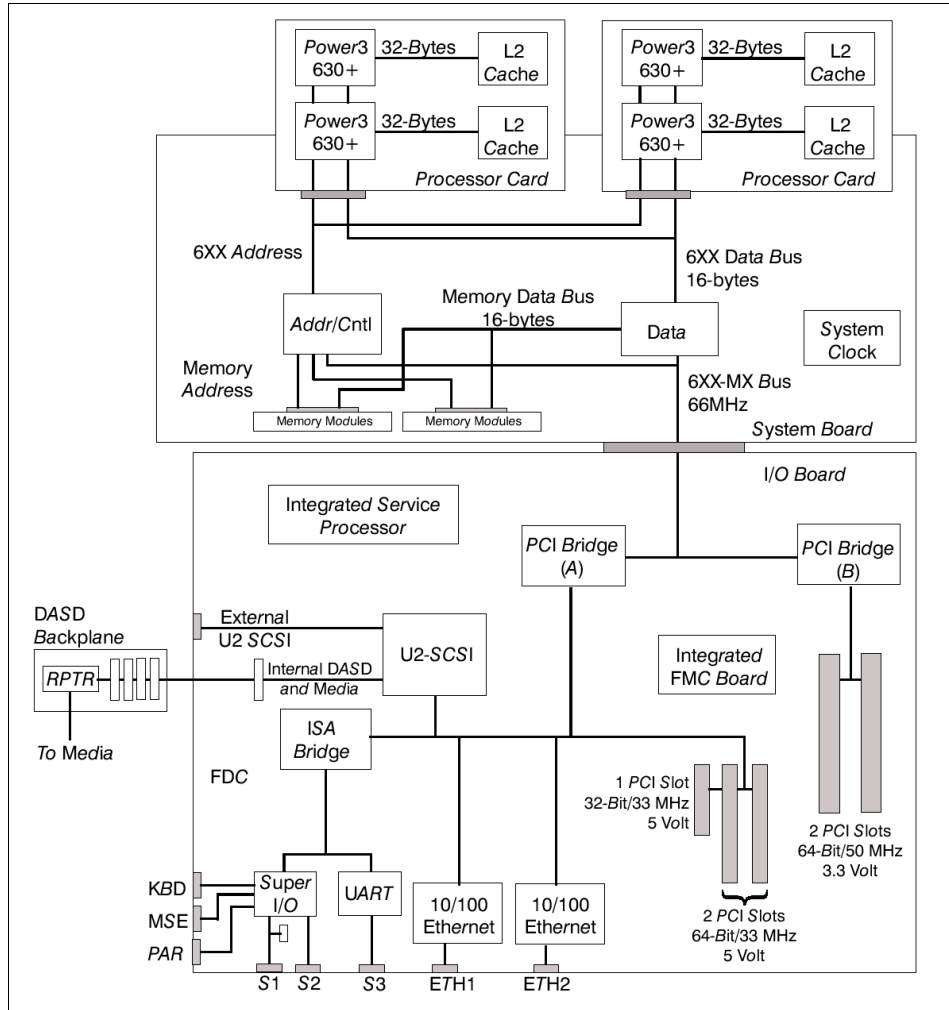


Figure 3-40 System architecture diagram for p640 Model B80

The processor subsystem is implemented on pluggable processor cards. The p640 Model B80 system has two processor card slots for support of 1- to 4-way POWER3-II processors running at either 375 MHz or 450 MHz. The processor card also contains the master oscillator for the CPU and memory sub-system.

The memory controller function is implemented in the chipset located on the system backplane. This chipset supports a 128-bit data path to memory along with an 8-bit ECC code that provides single bit correction, double-bit error detection, and 4-bit packet error detection. This chipset provides both an

interface to memory and a 6XX-MX bus for I/O. The system has two pluggable memory cards and support up to 16 GB ECC SDRAM memory.

The I/O subsystem is attached through the 6XX-MX bus provided by the chipset modules in the system backplane. There are two chip modules that provide a PCI bus for the I/O, as well as the Winbond ISA bridge, which provides an ISA bus for additional native I/O.

3.8.4 RAS features

The following are the RAS features of the p640 Model B80:

- ▶ Dynamic Processor Deallocation
- ▶ ECC L2 cache, SDRAM
- ▶ Service processor
- ▶ Hot-swappable disk bays
- ▶ Hot plug power supplies and cooling fans
- ▶ Redundant hot-plug cooling fans
- ▶ Optional redundant hot-plug power
- ▶ NEBS
- ▶ HACMP supported

3.8.5 Minimum and standard features

This section describes the system minimum and standard features for p640 Model B80

Table 3-28 summarizes the minimum and standard feature for p640 Model B80.

Table 3-28 Model B80 minimum and standard features

Features (min./standard)	Description
Microprocessor type	64-bit POWER3-II
Min. number of processors	1-way 375 MHz POWER3-II with 4 MB L2 cache
Memory (minimum)	256 MB
Memory bus width	Quad 512-bit
Data/Instruction (L1) cache	32 KB/64 KB
Level 2 (L2) cache	4 MB

Features (min./standard)	Description
Internal disk bays	Four; front accessible, hot swappable
Internal disk drive	36.4 GB Ultra SCSI
Media bays	One; optional front accessible tape/CD-ROM/DVD-RAM
PCI slots	Five PCI
PCI bus width	32- and 64-bit
Standard Ports	Three serial ports Keyboard and mouse ports
Integrated SCSI Adapters	Two integrated Ultra2 SCSI controllers (one internal and one external)
Integrated LAN adapter	10/100 Mbps Ethernet controller with two ports
System support (service) processor	Yes
Operating system version	AIX 5L Versions 5.1/5.2 Linux 2.4 available from one or more IBM Linux Distribution Partners
Warranty	One year

3.8.6 System expansion

This section describes the available features for the p640 Model B80 expansion.

Table 3-29 summarizes the possible maximum processor, memory, and storage expansion features for p640 Model B80.

Table 3-29 Model B80 system expansion features

System expansion features	Description
SMP configuration	2-, 3-, or 4-way 375 MHz POWER3-II (4 MB or 8 MB L2 cache/processor) 2- or 4-way 450 MHz POWER3-II (8 MB L2 cache/processor)
Memory	Up to 16 GB
Internal disk drive	Up to 293.6 GB (36.4 GB and 73.4 GB drives available)
External storage	IBM 2104 Expandable Storage Plus (Ultra3 SCSI), IBM 7133 Serial Disk System (SSA)

3.8.7 Features

Refer to Table 3-31 on page 150 for information about the features for p640 Model B80.

3.8.8 Configuration notes

The following section provides the major configuration notes for the p640 Model B80.

Processors

The following are the configuration notes regarding the processors:

- ▶ A 1-way (FC 4361) or 2-way (FC 4362) 375 MHz POWER3-II with 4 MB of L2 cache per processor cannot be used with the 2-way (FC 4365) 375 MHz POWER3-II with 8 MB of L2 cache per processor.

Memory

The following are the configuration notes regarding the memory:

- ▶ A maximum of 16 GB of system memory is allowed.
- ▶ A minimum of 256 MB of memory is required.
- ▶ One or two memory cards may be installed.
- ▶ The first memory card must be installed in memory card slot 1.
- ▶ 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) and then moving up (Figure 3-41 on page 146).
- ▶ All slots in the first memory card must be filled before a second memory card is installed.

Figure 3-41 on page 146 shows the memory card locations.




Figure 3-41 Memory card locations

Power supplies

The following are the configuration notes regarding the power supplies:

- ▶ AC or DC redundant power supply
- ▶ First power supply unit must be specified:
 - AC (FC 9172)
 - DC (FC 9175)
- ▶ Redundant power supply
 - AC (FC 6280)
 - DC (FC 6281)

Power cords

The following are the configuration notes regarding the power cords:

- ▶ Power cord specify feature codes 98xx provide a power cord to power an external I/O device from a wall-type outlet. These do not work with the rack Power Distribution Unit (PDU). To plug the I/O device into a 7014-S00, T00, or T42 rack PDU, order FC 6095.
- ▶ Power cord FC 9911 is the standard IBM rack power cord that goes from the B80 to the PDU. Customers will receive FC 9911 if they do not specify a power cord. Power cord FC 9900-9909 are available to customers who want to plug the B80 into a regular wall-type outlet.
- ▶ Refer to Appendix C, “Power cord features” on page 907 for further information.

Racks

The following are the configuration notes regarding the racks:

- ▶ The IBM 7014-T00 rack, which supports the pSeries 640 with DC power, provides eight power supplies.
- ▶ A maximum of seven pSeries 640 systems can be mounted in a T00 rack.
- ▶ A maximum of six pSeries 640 systems can be mounted in an S00 rack.
- ▶ A maximum of eight pSeries 640 systems can be mounted in a T42 rack.
- ▶ In order to mount more than four p640 systems with DC power in a T00 rack, FC 0986 must be specified so the configuration will receive special request services from the Center for Customized Solutions.
- ▶ DC Powered pSeries 640 Model B80 systems are supported only in the 7014-T00 rack equipped with DC power distribution panel (FC 6117). The 7014-T00 rack power distribution panel (FC 6117) provides attachment for up to four 3-pin to 3-pin PDP power cables (FC 6285) and up to four 3-pin to 5-pin PDP power cables (FC 6286). This will allow a maximum of up to four B80s with redundant power or up to seven B80s without redundant power. Up to four 7133 SSA drawers can also be installed with the four redundant power B80s in the same 7014-T00 rack. DC 7133 SSA drawers are shipped automatically with the correct power cables.
- ▶ FC 6117 is a feature of the 7014-T00 rack. FC 6117 provides the proper capacity circuit breakers for use with the pSeries 640 Model B80 system based on the inclusion of rack FC 0178 (Rack Content Specify, 7026-B80) on the 7014-T00 rack order.

Disks, media, and boot devices

The following are the configuration notes regarding the disks, media, and boot devices:

- ▶ Media bay 1 (refer to Figure 3-42 on page 148 and Table 3-30 on page 148) can accommodate a CD-ROM, DVD-RAM, tape drive, or other media device.
- ▶ All DASD and media devices are driven by the internal/integrated SCSI port only.

Figure 3-42 on page 148 shows the disks and media bay locations.




Figure 3-42 Disks and media bay locations

Table 3-30 describes the disks and media bay locations and the SCSI IDs.

Table 3-30 Disk and media bay locations and SCSI-IDs

Item number	Bay location	Description	SCSI-ID
1	Bay D01	Media bay	SCSI ID 0
2	Bay D02	Disk bay	SCSI ID 8
3	Bay D03	Disk bay	SCSI ID 9
4	Bay D04	Disk bay	SCSI ID 10
5	Bay D05	Disk bay	SCSI ID 11

Note: The SCSI bus IDs listed are the recommended values. The SCSI IDs shown for media devices indicate how the devices are set when shipped from the factory. Field installations might not comply with these recommendations.

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ Slots 1 and 2 are 64-bit and run at 50 MHz.
- ▶ Slots 3 and 4 are 64-bit and run at 33 MHz.
- ▶ Slot 5 is 32-bit and runs at 33 MHz.
- ▶ 32-bit slot is half-size.
- ▶ 64-bit slots are full-size.

- ▶ The PCI Advanced SerialRAID Plus adapter (FC 6230) can be selected with or without the Fast-Write Cache Option card (FC 6235).
- ▶ The SSA Fast-Write Cache Option card (FC 6235) is a cache upgrade for the SSA Advanced SerialRAID Plus adapter (FC 6230) and can only be selected along with that adapter (maximum of 1 per FC 6230).
- ▶ The Gigabit Fibre Channel Adapter (FC 6227) is restricted to only one adapter in slots 1 or 2 and one adapter in slots 3, 4, or 5.
- ▶ Refer to Appendix B, “Adapter placement guidelines” on page 785 for additional I/O adapter configuration notes.

Graphics adapters

The following are the configuration notes regarding the graphics adapters:

- ▶ Graphics adapter, keyboard, and mouse are not required in the minimum configuration.
- ▶ The maximum number of graphics adapters supported in the pSeries 640 is one.

Hot-plug options

The following are the configuration notes regarding the Hot-Plug options:

- ▶ It is not necessary to power down the system to install/remove/replace hot-plug options. Refer to p640 Model B80 product documentation for further information.
- ▶ The following options are hot-pluggable:
 - Hot-plug power supplies (if the drawer has two power supplies and only one needs to be removed)
 - Hot-plug fans

3.9 Features for entry-level systems

This section lists the internal features that can be added to the entry-level systems configuration. The status of a feature is indicative of these qualifications:

X	Available. Indicates features that are available and are orderable on the specified model.
R	Must be removed. Indicates this feature is not supported and must be removed during a model conversion.
W	Withdrawn. Indicates a feature that is no longer available to order.

Features not listed in the provided categories, indicate that the feature is not supported on this model. Some categories, such as keyboards, mice, language specified codes, and power cords are not included.

Table 3-31 reports the features available for the 7026 Model B80, 7028 Model 6C1, 7028 Model 6E1, 7028 Model 6C4, 7028 Model 6E4, 7029 Model 6C3, 7029 Model 6E3, 7043 Model 150, 7044 Model 170, 9112 Model 265, and 9114 Model 275.

Table 3-31 Features for entry-level systems

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
Host attachment adapters												
2751	S/390® ESCON Channel PCI Adapter		X	X								
8397	SP Switch2 PCI Attachment Adapter				X							
8398	SP Switch2 PCI-X Attachment Adapter				X							
Asynchronous adapters												
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI	X	X	X	X	X	X	X	X	X	X	X
2944	128-Port Asynchronous Controller, PCI bus	X	X	X	X	X	X	X	X	X	X	X
SCSI adapters												
6204	PCI Universal Differential Ultra SCSI Adapter	X	X	X	X	X	X	X		X	X	X
6207	PCI Differential Ultra SCSI Adapter								X			
6206	PCI Single-Ended Ultra SCSI Adapter	X	X	X					X	X	X	
6205	PCI Dual Channel Ultra2 SCSI Adapter	W							W	W		
2494	PCI 3-Channel Ultra2 SCSI RAID Adapter								W	W		
6203	PCI Dual Channel Ultra3 SCSI Adapter	X	X	X	X	X	X	X	X	X	X	X
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter	X	X	X	X	X	X	X	X	X	X	X
5709	ULTRA320 SCSI RAID, Daughter Card						X	X				X
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter		X	X	X	X	X	X			X	X
5712	PCI-X Dual Channel Ultra320 SCSI Adapter		X	X	X	X	X	X		X	X	X
Serial adapters												
6231	128 MB DRAM Option Card	X	X	X	X	X	X	X	X	X	X	X

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
6225	IBM Advanced SerialRAID Adapter								W			
6230	Advanced SerialRAID Plus Adapter	X	X	X	X	X	X	X	X	X	X	X
6235	32 MB Fast-Write Cache Option Card	X	X	X	X	X	X	X	X	X	X	X
6215	PCI SSA Multi-Initiator/RAID EL Adapter								W			
6222	SSA Fast-Write Cache								W			
Fibre adapters												
6227	Gigabit Fibre Channel PCI Adapter	W								W		
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI	X	X	X	W	W				X	W	
6239	2 Gigabit Fibre Channel PCI-X Adapter	X	X	X	X	X	X	X		X	X	X
Graphics adapters												
2848	POWER GXT135P Graphics Accelerator	X	X	X	X	X			X	X	X	
2849	POWER GXT135P Graphics Accelerator with Digital Support	X	X	X	X	X	X	X	X	X	X	X
2830	POWER GXT130P Graphics Accelerator	W	W	W					W	W		
2845	POWER GXT550P Graphics Adapter/C - PCI								W			
2823	POWER GXT2000P Graphics Accelerator								W	W		
2842	POWER GXT4500P Graphics Accelerator	X			X	X			X	X	X	X
2843	POWER GXT6500P Graphics Accelerator	X			X	X				X	X	X
2838	POWER GXT120P Graphics Adapter - PCI								W			
2825	POWER GXT3000P Graphics Accelerator								W	W		
2851	POWER GXT250P Graphics Adapter - PCI								W			
2852	POWER GXT255P Graphics Adapter - PCI								W			
2841	POWER GXT300P Graphics Accelerator								W	W		
2826	POWER GXT4000P Graphics Accelerator									W		
2827	POWER GXT6000P Graphics Accelerator									W		

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
LAN adapters												
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	X	X	X	X	X	X	X	X	X	X	X
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	X	X	X	X	X	X	X	X	X	X	X
4951	IBM Four Port 10/100 Ethernet Adapter	W							X	W		
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	X	X	X	X	X	X	X		X	X	
4959	IBM Token-Ring PCI Adapter	X	X	X	X	X	X	X	X	X	X	X
2920	IBM PCI Token-Ring Adapter								W			
2946	Turboways 622 Mbps PCI MMF ATM Adapter	X	X	X	X	X	X	X		X	X	X
2998	TURBOWAYS® 25 ATM PCI Adapter								W	W		
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	X	X	X	X	X	X	X		X	X	X
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	X	X	X	X	X	X	X		X	X	X
2985	PCI Ethernet BNC/RJ-45 Adapter									W		
2987	PCI Ethernet AUI/RJ-45 Adapter									W		
2968	IBM 10/100 Mbps Ethernet PCI Adapter	W							W	W		
4962	10/100 Mbps Ethernet PCI Adapter II	X	X	X	X	X	X	X	X	X	X	X
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	W			W	W			W	W		
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	W			W	W			W	W		
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	W							W	W		
2969	Gigabit Ethernet - SX PCI Adapter	X	X	X	X	X				X	X	
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	X	X	X	X	X	X	X		X	X	X
2975	10/100/1000 Base-T Ethernet PCI Adapter	X	X	X	X	X				X	X	
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	X	X	X	X	X	X	X		X	X	X
2988	TURBOWAYS 155 PCI MMF ATM Adapter	W							W	W		
2963	TURBOWAYS 155 PCI UTP ATM Adapter	W							W	W		

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
WAN adapters												
2962	2-Port Multiprotocol PCI Adapter	X	X	X	X	X	X	X	X	X	X	X
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	X	X	X	X	X	X	X	X	X		X
6310	IBM ARTIC960RxD Quad Digital Trunk PCI Adapter	X	X	X	X	X	X	X	X	X		X
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter								W	W		
6312	Quad Digital Trunk Telephony PCI Adapter	X	X	X	X	X	X	X	X	X	X	X
Miscellaneous adapters												
4958	PCI Cryptographic Coprocessor	W							W	W		
6311	IBM ARTIC960RxF PCI Adapter	W							W	W		
8244	Audio PCI Adapter for Workstations					X	X	X			X	X
4963	PCI Cryptographic Coprocessor (FIPS-4)	X	X	X	X	X			X	X		
4960	IBM e-business Cryptographic Accelerator	X	X	X	X	X	X	X	X	X	X	X
2733	IBM Long-wave Serial HIPPI PCI Adapter		W	W	W	W						
2732	IBM Short-wave Serial HIPPI PCI Adapter		X	X	X	X	X					
2737	Keyboard/Mouse Attachment Card - PCI				X	X						
2639	Ultimedia Video Capture Adapter/S - PCI								W	W		
Asynchronous cables												
8135	64-Port to 128-Port Pin-Out Converter	W							W	W		
2936	Asynchronous Cable EIA-232/V.24	X	X	X	X	X	X	X	X	X	X	X
2945	Asynchronous Terminal Cable - EIA-422A	W							W	W		
2934	Asynchronous Terminal/Printer Cable EIA-232	X	X	X	X	X	X	X	X	X	X	X
3926	Asynch Printer/Terminal Cable, 9-pin to 25-pin, 4M	X	X	X	X	X	X	X		X	X	X
8131	128-Port Asynchronous Controller Cable, 4.5 Meter	X	X	X	X	X	X	X	X	X	X	X
8132	128-Port Asynchronous Controller Cable, 23 cm (9-In.)	X	X	X	X	X	X	X	X	X	X	X
8121	Attachment Cable, Hardware Management Console to Host, 15-Meter				X							

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8120	Attachment Cable, Hardware Management Console to Host, 6-Meter				X							
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232	X	X	X	X	X	X	X	X	X	X	X
8138	Enhanced Remote Asynchronous Node 16-Port RS-422	W							W	W		
8133	RJ-45 to DB-25 Converter Cable	X	X	X	X	X	X	X	X	X	X	X
8136	Rack Mountable Remote Asynchronous Node 16-Port EIA-232	X	X	X	X	X	X	X			X	X
3925	Serial Port Converter Cable, 9-Pin to 25-Pin	X	X	X	X	X	X	X		X	X	X
Internal IDE cables												
4246	IDE 2-Drop Connector Cable		X	X								X
Internal SCSI cables												
4248	SCSI Connector Cable and Repeater Card		X	X								X
4259	SCSI 3-Drop Connector Cable		X	X								X
4249	SCSI 3-Drop Connector Cable		X	X								X
4250	Connector Cable, 6-slot PCI Riser to Media Bays				X	X						
4260	SCSI 2-drop Connector Cable and Terminator				X	X						
4247	SCSI 2-Drop Connector Cable		X	X								X
2445	Internal Ultra SCSI Cable Assembly (from PCI Single-Ended to Single-Ended Ultra SCSI devices)								W			
4254	SCSI Connector Cable				X	X						
4267	PCI to Ultra320 backplane cable						X	X				
4258	SCSI Cable, PCI Adapter to Hot-Swap Disks		X	X								X
4263	SCSI Cables (Power and Logic), PCI Riser to SCSI LVD Media Device; Mounting Tray						X	X				X
4266	SCSI Cables (Power and Logic), PCI Riser to SCSI Media Device; Mounting Tray						X	X				X

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
External SCSI cables												
2424	0.6M 16-bit SCSI-2 System to System Cable	X	X	X	X	X	X	X	X	X	X	X
2425	2.5M 16-bit SCSI-2 System to System Cable	X	X	X	X	X	X	X	X	X	X	X
2116	PCI SCSI Adapter To 2-Port, 16-Bit Differential Ext Device Cable	W	W	W					W	W		
2112	PCI SCSI Adapter To 2-Port, 8-Bit Differential Ext Device Cable	W							W	W		
2115	PCI SCSI Adapter To 2-Port, 16-Bit SE External Device Cable	W	W	W					W	W		
2111	PCI SCSI Adapter To 2-Port, 8-Bit SE External Device Cable	W							W	W		
2113	PCI SCSI Adapter To Single Port 8-Bit SE External Device Cable	W							W	W		
2118	Converter Cable, VHDCI to P, Mini-68 pin to 68 pin, 0.3M	X	X	X	X	X	X	X	X	X	X	X
2114	PCI SCSI Adapter 16-Bit Differential External Y Cable	X	X	X	X	X	X	X	X	X	X	X
2117	PCI SCSI Adapter 16-bit SE External Y Cable	W							W	W		
Fibre cables												
2456	LC-SC 50 Micron Fibre Converter Cable	X	X	X	X	X	X	X		X	X	X
2459	LC-SC 62.5 Micron Fibre Converter Cable	X	X	X	X	X	X	X		X	X	X
Cables for graphic adapters												
4241	13W3 to 13W3 DDC/ID Bits Switchable Display Cable								W			
4240	DDC 13W3 to 13W3 Display Cable								W			
4213	13W3 to 15-Pin D-Shell Converter Cable								W			
4219	6091-19i (77Hz-ISO), POWERdisplay 17, 19, 20 Monitor Cable								W			
4237	13W3 to 15-pin DDC/ID Bits Switchable Display Cable	W							W	W		

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
4235	15 Pin-D Shell to 13W3 Display Cable								W			
4238	DDC 15 Pin to 13W3 Display Cable	W							W	W		
4239	15-D Adapter to 3BNC ID Cable	W							W	W		
4217	6091 Attachment Cable	W							W	W		
4242	6 Foot Extender Cable for Displays (15 pin D-shell to 15 pin D-shell)	X	X		X	X	X	X	X	X		X
Transceivers for LAN adapters												
4223	Ethernet 10Base2 Transceiver								X	X		
4224	Ethernet 10BaseT Transceiver								W	W		
Cables for WAN adapters												
2861	ARTIC960Hx 4-Port EIA-232 Cable	X	X	X	X	X	X	X	X	X		X
2865	ARTIC960Hx 4-Port EIA-530 Cable	W							W	W		
2862	ARTIC960Hx 4-Port RS-449 Cable	W							W	W		
2864	ARTIC960Hx 4-Port V.35 (DTE) Cable	X	X	X	X	X	X	X	X	X		X
2863	ARTIC960Hx 4-Port X.21 Cable	X	X	X	X	X	X	X	X	X		X
2710	ARTIC960Hx/RxD 4-Port E1 RJ-45 Cable	X	X	X	X	X	X	X	X	X		X
2709	ARTIC960Hx/RxD 4-Port T1 RJ-45 Cable	X	X	X	X	X	X	X	X	X		X
2873	IBM ARTIC960RxD Quad DTA, E1, 120 Ohm Balanced, 3M 4-Port Cable	X	X	X	X	X	X	X	X	X		
2874	IBM ARTIC960RxD Quad DTA, E1, 120 Ohm Balanced, 7.5M Extension Cable	X	X	X	X	X	X	X	X	X		
2875	IBM ARTIC960RxD Quad DTA, E1, 75 Ohm Unbalanced-Grounded, 1.8M 4-Port Cable	W	W	W					W	W		
2876	IBM ARTIC960RxD Quad DTA, E1, 75 Ohm Unbalanced-Ungrounded, 1.8M 4-Port Cable	W							W	W		
2871	IBM ARTIC960RxD Quad DTA, T1, 100 Ohm, 3M 4-Port Cable	X	X	X	X	X	X	X	X	X		
2872	IBM ARTIC960RxD Quad DTA, T1, 100 Ohm, 15M Extension Cable	X	X	X	X	X	X	X	X	X		

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
2951	Cable, V.24/EIA-232	X	X	X	X	X	X	X	X	X	X	X
2952	Cable, V.35	X	X	X	X	X	X	X	X	X	X	X
2953	Cable, V.36/EIA-499	X	X	X	X	X	X	X	X	X	X	X
2954	Cable, X.21	X	X	X	X	X	X	X	X	X	X	X
2877	IBM ARTIC960RxD Quad DTA, H.100, 4-Drop Cable	X	X	X	X	X	X	X	X	X		X
4353	H-100 Bus 8-position Cable		X	X	X	X	X	X				X
2878	IBM ARTIC960RxD Quad DTA, SC Bus, 5-Drop Cable									W		
2879	IBM ARTIC960RxD Quad DTA, Four H.100 to SC bus Converter Cable									W		
Cables for racks												
4692	Rack Status Beacon Cable, Junction Box Daisy Chain				X	X						
4691	Rack Status Beacon Cable, Junction Box To Drawer Or Status Beacon				X	X						
3124	Serial to Serial Port Cable for Drawer/Drawer	X	X	X	X	X	X	X	X	X	X	X
3125	Serial to Serial Port Cable for Rack/Rack	X	X	X	X	X	X	X	X	X	X	X
3254	RIO Cable Support Brackets				X							
6001	Power Control Cable (SPCN) - 2 meter				X							
6006	Power Control Cable (SPCN) - 3 meter				X							
6008	Power Control Cable (SPCN) - 6 meter				X							
6007	Power Control Cable (SPCN) - 15 meter				X							
3147	RIO-2 (Remote I/O-2) Cable, 3.5M				X							
3148	RIO-2 (Remote I/O-2) Cable, 10M				X							
SCSI disk drives												
2900	4.5 GB Ultra-SCSI 16-bit Disk Drive								W			
3027	9.1 GB 10,000 RPM Ultra-SCSI Disk Drive								W	W		

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
2908	9.1 GB Ultra-SCSI 16-bit 1-in. (25mm) High Disk Drive								W	W		
3025	9.1 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly	W										
3102	18.2 GB 10,000 RPM Ultra SCSI Disk Drive		W	W					W	W	W	
2909	18.2 GB Ultra-SCSI 1" High (25mm) Disk Drive								W	W		
3026	18.2 GB 10,000 RPM Ultra2 SCSI Disk Drive Assembly	W										
3263	18.2 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly		W	W							W	
3157	18.2 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly				W	W						
3119	36.4 GB 10,000 RPM Ultra SCSI Disk Drive		X	X					X	X	X	
3264	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly		X	X							X	
3129	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	X										
3158	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly				X	X						
3273	36.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly						X	X				X
3280	36.4 GB 15,000 RPM Ultra3 SCSI Disk Drive Assembly		X	X							X	
3277	36.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly	X			X	X	X	X				X
3118	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive		X	X					X	X		
3265	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly		X	X							X	
3159	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	X			X	X						
3274	73.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly						X	X				X

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
3281	73.4 GB 15,000 RPM Ultra3 SCSI Disk Drive Assembly		X	X							X	
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly				X	X	X	X				X
3276	146.8 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly		X	X							X	
3275	146.8 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly				X	X	X	X				X
SCSI mounting hardware												
6509	Media Bay Disk Drive Mounting Kit								X			
6560	Media Bay Disk Mounting Kit									W		
6566	Mounting Hardware for Disk in Media Bay		X	X							X	
6561	Media Bay Cooling Fan and Drive Mounting Kit									X		
Flat panel and graphic monitors												
3636	L200p Flat Panel Monitor		X	X	X	X	X	X	X	X	X	X
3635	T210 Flat-Panel Monitor		W	W	W	W			W	W	W	
3626	IBM P202 Color Monitor, Business Black, and Cable								W			
3628	IBM P260/P275 Color Monitor, Business Black, and Cable	X	X	X	X	X	X	X	X	X	X	X
3623	IBM P72 Color Monitor, Business Black, Captured Cable								W			
3627	IBM P76/P77 Color Monitor, Business Black, Captured Cable	W	W	W	W	W			W	W	W	
3625	IBM P92 Color Monitor, Business Black, and Cable								W			
3637	IBM T541H 15" TFT Color Monitor, Business Black, Captured Cable	X	X	X	X	X	X	X	X	X	X	X
3622	IBM P202 Color Monitor								W			
3630	IBM P260/P275 Color Monitor, Pearl White, and Cable	W							W	W		
3620	IBM P72 Color Monitor								W			

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
3629	IBM P76/P77 Color Monitor, Pearl White, Captured Cable	W							W	W		
3621	IBM P92 Color Monitor								W			
Memory DIMMs												
4149	64 MB SDRAM DIMM, 168 Pin								W			
4150	128 MB SDRAM DIMM, 168 Pin								X			
4169	256 MB SDRAM DIMM, 168 Pin								X			
4110	256 MB(2X128 MB) DIMMS, 200 pin 10NS SDRAM	X								X		
4119	512 MB (2X256 MB) DIMMs, 200-pin 10NS SDRAM									W		
4120	512 MB (2x256 MB) DIMMs, 200-pin 10NS SDRAM	X	X	X						X	X	
4121	1024 MB (2 X 512 MB) DIMMs, 200-pin 10NS SDRAM	X	X	X						X	X	
4451	1024 MB (4x256 MB) DIMMs, 208-pin, 8NS DDR SDRAM				X	X						
4446	1024 MB (4x256 MB) DIMMs, 208-pin, 8NS DDR SDRAM						X	X				X
4452	2048 MB (4x512 MB) DIMMs, 208-pin, 8NS DDR SDRAM				X	X						
4447	2948 MB (4x512 MB) DIMMs, 208-pin, 8NS DDR SDRAM						X	X				X
4453	4096 MB (4x1024 MB) DIMMs, 208-pin, 8NS Stacked DDR SDRAM				X	X						
4448	4096 MB (4x1024 MB) DIMMs, 208-pin, 8NS DDR SDRAM						X	X				X
4454	8192 MB (4x2048 MB) DIMMs, 208-pin, 8NS Stacked DDR SDRAM				X	X						
4449	8192 MB (4x2048 MB) DIMMs, 208-pin, 8NS DDR SDRAM						X	X				X
8093	1024 MB (2x512 MB) DIMMs, Express Configurations, Factory Only		X	X								

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8094	1024 MB (2x512 MB) DIMMs, Express Configurations, Factory Only		X	X								
8095	1024 MB (2x512 MB) DIMMs, Express Configurations, Factory Only		X	X								
8156	1024 MB (4x256 MB) DIMMs, Express Configurations, Factory Only						X	X				
8157	2048 MB (4x512 MB) DIMMs, Express Configurations, Factory Only						X	X				X
8158	4096 MB (4x1024 MB) DIMMs, Express Configurations, Factory Only						X	X				X
8159	8192 MB (4x2048 MB) DIMMs, Express Configurations, Factory Only						X	X				
8081	1024 MB (4x256 MB) DIMMs, Express Configurations, Factory Only				X	X						
8080	2048 MB (4x512 MB) DIMMs, Express Configurations, Factory Only				X	X						
8079	4096 MB (4x1024 MB) DIMMs, Express Configurations, Factory Only				X	X						
Memory expansion card												
4098	Memory Expansion Feature (16 POS)	X										
Power supplies												
6275	Power Supply, 250 Watt AC, Hot-swap, Redundant		X	X								X
6277	Power Supply, 250 Watt AC, Hot-swap		X	X								X
6273	Power Supply, 645 Watt AC, Hot-swap, Base and Redundant				X	X						
6266	Power Supply, 680 Watt AC, Hot-swap, Base and Redundant						X	X				X
9172	Power Specify, AC	X			X							
6274	Power Supply, 645 Watt DC, Hot-swap, Base and Redundant				X							
9175	Power Specify, DC	X			X							

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
6285	DC power cable - 3 pin	X			X							
6286	DC power cable - 5 pin	X			X							
6280	Redundant AC Power Supply, 560 Watt	X										
6281	Redundant DC Power Supply, 544 Watt	X										
Diskette drives												
2605	Diskette Drive (Black bezel)				X	X						
2611	Slimline Diskette Drive (Black bezel)						X	X				X
SCSI CD-ROM/DVD-RAM												
2624	32X (Max) SCSI-2 CD-ROM Drive	X							X	X		
2623	4.7 GB SCSI-2 DVD-RAM Drive (Black Bezel)	X	X	X	X	X	X	X	X	X	X	X
2627	4.7 GB SCSI-2 DVD-RAM Drive (White Bezel)	W							W	W		
SCSI tapes												
6134	60/150 GB 16-bit 8mm Internal Tape Drive	X	X	X	X	X	X	X	X	X	X	X
6159	12/24 GB 4mm Internal Tape Drive								W	W		
6158	20/40 GB 4mm Internal Tape Drive	X	X	X	X	X	X	X	X	X	X	X
6156	20/40 GB 16-bit 8mm Internal Tape Drive (Black Bezel)		W	W	W	W					W	W
6154	20/40 GB 16-bit 8mm Internal Tape Drive (White Bezel)									W		
6120	IBM 80/160 GB Internal Tape Drive with VXA Technology	X	X	X	X	X	X	X	X	X	X	X
IDE CD-ROM/DVD-ROM												
2634	16X/48X(max) IDE DVD-ROM Drive		X	X	X	X						X
2633	IDE CD-ROM Drive (Black bezel)		X	X	X	X						X
2640	IDE Slimline DVD-ROM Drive						X	X				X
LPAR enablement												
9556	6 Slot PCI Riser (Initial order only)				X	X						

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
6556	6 Slot PCI Riser (MES order only)				X	X						
6575	Enhanced Planar (MES order only)				X	X						
9581	RIO-2 Enablement Indicator, Factory Only				X							
9575	RIO Ports and LPAR Enablement (Initial order only)				W	W						
6576	LPAR Enablement (MES order only)				X	X						
Multimedia/graphic devices												
6308	Personal Microphone								W			
8422	Spaceball 3D Input Device					X			X	X	X	X
8741	3-Button Mouse - Business Black	X	X	X	X	X	X	X	X	X	X	X
6041	3-Button Mouse	X							X	X		
8841	Mouse - Business Black with Keyboard Attachment Cable				X	X	X	X				
8679	IBM Multimedia Kit for RS/6000 (US, Canada)								W			
3753	Options Library (hard copy)								W	W		
8423	Magellan XT 3D Input Device					X			X	X	X	X
Miscellaneous												
6166	Acoustic Package for Deskside Systems			X								
6557	Redundant Cooling				X	X						
3752	Service Package	X	X	X					X	X	X	
8227	Tie-down Strap Security Device								X	X		
Express configurations												
8167	Feature Code Indicator for Express Configuration -- pSeries 615(110C or 110E)						X	X				
8169	Feature Code Indicator for Express Configuration -- pSeries 615 (120C or 120E)						X	X				
8170	Feature Code Indicator for Express Configuration -- pSeries 615 (140C or 140E)						X	X				

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8171	Feature Code Indicator for Express Configuration -- pSeries 615 (180C or 180E)						X	X				
8172	Feature Code Indicator for Express Configuration -- pSeries 615 (220C or 220E)						X	X				
8173	Feature Code Indicator for Express Configuration -- pSeries 615 (240C or 240E)						X	X				
8175	Feature Code Indicator for Express Configuration -- pSeries 615 (280C or 280E)						X	X				
8188	Feature Code Indicator for Express Configuration -- pSeries 615(224C or 224E)						X	X				
8189	Feature Code Indicator for Express Configuration -- pSeries 615(244C or 244E)						X	X				
8190	Feature Code Indicator for Express Configuration -- pSeries 615(284C or 284E)						X	X				
8160	Feature Code Indicator for Express Configuration -- pSeries 615X (111C or 111E)						X	X				
8161	Feature Code Indicator for Express Configuration -- pSeries 615X (121C or 121E)						X	X				
8162	Feature Code Indicator for Express Configuration -- pSeries 615X (141C or 141E)						X	X				
8163	Feature Code Indicator for Express Configuration -- pSeries 615X (181C or 181E)						X	X				
8164	Feature Code Indicator for Express Configuration -- pSeries 615X (221C or 221E)						X	X				
8165	Feature Code Indicator for Express Configuration -- pSeries 615X (241C or 241E)						X	X				
8166	Feature Code Indicator for Express Configuration -- pSeries 615X (281C or 281E)						X	X				
8191	Feature Code Indicator for Express Configuration -- pSeries 615X (225C or 225E)						X	X				
8192	Feature Code Indicator for Express Configuration -- pSeries 615X (245C or 245E)						X	X				

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8193	Feature Code Indicator for Express Configuration -- pSeries 615X (285C or 285E)						X	X				
8034	Feature Code Indicator for Express Configuration -- pSeries 630X (121C or 121E)				X	X						
8179	Feature Code Indicator for Express Configuration -- pSeries 630X (221C or 221E)				X	X						
8036	Feature Code Indicator for Express Configuration -- pSeries 630X (223C or 223E)				X	X						
8038	Feature Code Indicator for Express Configuration -- pSeries 630X (421C or 421E)				X	X						
8096	Express Configuration -- pSeries 610 (100E/100C)		X	X								
8097	Express Configuration -- pSeries 610 (200E/200C)		X	X								
8098	Express Configuration -- pSeries 610 (250E/250C)		X	X								
8088	Express Configuration -- pSeries 630 (100C)				W							
8103	Feature Code Indicator for Express Configuration -- pSeries 630X (101C)				W							
8089	Express Configuration -- pSeries 630 (200C)				W							
8104	Feature Code Indicator for Express Configuration -- pSeries 630X (201C)				W							
8090	Express Configuration -- pSeries 630 (400C)				W							
8105	Feature Code Indicator for Express Configuration -- pSeries 630X (401C)				W							
8033	Feature Code Indicator for Express Configuration -- pSeries 630 (120C or 120E)				X	X						
8177	Feature Code Indicator for Express Configuration -- pSeries 630(220C or 220E)				X	X						
8110	Feature Code Indicator for Express Configuration -- pSeries 630 (140C or 140E)				X	X						
8112	Feature Code Indicator for Express Configuration -- pSeries 630 (240C or 240E)				X	X						

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8035	Feature Code Indicator for Express Configuration -- pSeries 630 (222C or 222E)				X	X						
8037	Feature Code Indicator for Express Configuration -- pSeries 630 (420C or 420E)				X	X						
8114	Feature Code Indicator for Express Configuration -- pSeries 630 (440C or 440E)				X	X						
8085	Express Configuration -- pSeries 630 (100E)					W						
8100	Feature Code Indicator for Express Configuration -- pSeries 630X (101E)					W						
8086	Express Configuration -- pSeries 630 (200E)					W						
8101	Feature Code Indicator for Express Configuration -- pSeries 630X (201E)					W						
8087	Express Configuration -- pSeries 630 (400E)					W						
8102	Feature Code Indicator for Express Configuration -- pSeries 630X (401E)					W						
Disk bays												
6568	Ultra3 SCSI Backplane for Hot-swap Disks (4-pack)				X	X						
6567	Ultra3 SCSI Backplane for Hot-swap Disks		X	X							X	
6574	Ultra320 SCSI 4-Pack						X	X				X
6584	Ultra320 SCSI 4-Pack for Disk Mirroring						X	X				
Software preload												
7305	AAP Software Pre-Install Indicator	X	X	X	X	X	X	X	X	X	X	X
5005	Software Preinstall	X	X	X	X	X	X	X	X	X	X	X
Routings and indicators												
0986	CCS Customer Service Specify (US)	X	X	X	X	X	X	X	X	X	X	
0008	Customer Install MES Indicator (WT)	X										
5001	Customer Service Specify	X	X	X	X	X	X	X	X	X	X	
0700	HA Solution Indicator	X										

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
0704	HA Solution Indicator	W	W									
Processor cards												
5322	1-Way 333 MHz POWER3-II Processor Card		X	X								
5300	1-Way 375 MHz POWER3-II Processor Card		X	X								
5323	1-Way 450 MHz POWER3-II Processor Card										X	
5309	1-way 450 MHz POWER3-II Processor Card		X	X								
5133	1-way 1.2 GHz POWER4+ Processor Card				X	X						
5126	1-way 1.45 GHz POWER4+ Processor Card				X	X						
5223	1-way 1.45 GHz POWER4+ Processor Card											X
5131	1-way 1.0 GHz POWER4 Processor Card				W	W						
5218	1-way 1.0 GHz POWER4 Processor Card											X
5134	2-way 1.2 GHz POWER4+ Processor Card				X	X						
5127	2-way 1.45 GHz POWER4+ Processor Card				X	X						
5132	2-way 1.0 GHz POWER4 Processor Card				W	W						
5224	2-way 1.45 GHz POWER4 Processor Card											X
5234	2-way 1.45 GHz POWER4 Processor Card						X	X				
8108	1-way 1.45 GHz POWER4+ Processor Card, Express Configurations, Factory Only				X	X						
8109	2-way 1.45 GHz POWER4+ Processor Card, Express Configurations, Factory Only				X	X						
8182	1-way 1.45 GHz POWER4+ Processor Card, Express Configurations, Factory Only											X
8184	2-way 1.45 GHz POWER4+ Processor Card, Express Configurations, Factory Only											X
8082	1-way 1.0 GHz POWER4 Processor Card, Express Configurations, Factory Only				W	W						
8083	2-way 1.0 GHz POWER4 Processor Card, Express Configurations, Factory Only				W	W						

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
8084	2-way 1.0 GHz POWER4 Processor Card, Express Configurations, Factory Only				W	W						
8106	1-way 1.2 GHz POWER4+ Processor Card, Express Configurations, Factory Only				X	X						
8107	2-way 1.2 GHz POWER4+ Processor Card, Express Configurations, Factory Only				X	X						
8148	1-way 1.2 GHz POWER4+ Processor Card, Express Configurations, Factory Only						X	X				
8149	2-way 1.2 GHz POWER4+ Processor Card, Express Configurations, Factory Only						X	X				
8187	2-way 1.45 GHz POWER4+ Processor Card, Express Configurations, Factory Only						X	X				
8091	1-way 450 MHz POWER3-II Processor Card, Express Configurations, Factory Only		X	X								
8092	1-way 450 MHz POWER3-II Processor Card, Express Configurations, Factory Only		X	X								
4363	PowerPC 604e 250 MHz Processor								X			
4348	PowerPC 604e 375 MHz Processor								X			
4349	1-Way 333 MHz POWER3-II Processor Card w/1 MB L2 Cache									X		
4361	1-Way 375 MHz POWER3-II Processor Card	X										
4360	1-Way 400 MHz POWER3-II Processor Card w/4 MB L2 Cache									X		
4364	1-Way 450 MHz POWER3-II Processor Card w/8 MB L2 Cache									X		
4362	2-Way 375 MHz POWER3-II Processor Card	X										
4365	2-WAY 375 MHz POWER3-II Processor Card (8 MB L2/Processor)	X										
4366	2-WAY 450 MHz POWER3-II Processor Card (8 MB L2/Processor)	X										
0514	2-Way 375 MHz POWER3-II Processor, Infrastructure Solutions, Factory Only	W										

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
0517	2-Way 450 MHz POWER3-II Processor, Infrastructure Solutions, Factory Only	W										
0516	1-Way 450 MHz POWER3-II Processor, Infrastructure Solutions, Factory Only		W									
5220	1-way 1.2 GHz POWER4+ Processor Card on Planar					X	X					
5222	2-way 1.2 GHz POWER4+ Processor Card on Planar					X	X					
5234	2-way 1.45 GHz POWER4+ Processor Card on Planar					X	X					
Rack PDU power cables												
4693	Rack Status Beacon Junction Box				X	X						
6285	Power cable, DC, Drawer (3-pin) to PDP (3-pin), 4M	X			X							
6286	Power Cable, DC, Drawer (3-pin) to PDP (5-pin)	X			X							
9900	Power Cord (4M) Specify -- United States/Canada	X	X		X	X						
9902	Power Cord (4M) Specify -- Denmark (250V, 10A)	X	X		X	X						
9905	Power Cord (4M) Specify -- Switzerland (250V, 10A)	X	X		X	X						
9908	Power Cord (4M) Specify -- Australia/New Zealand/Argentina (250V, 10A)	X	X		X	X						
9903	Power Cord (4M) Specify -- U.K. and Others (250V, 13A)	X	X		X	X						
9909	Power Cord (4M) Specify -- Thailand (250V, 15A)	X	X		X	X						
9901	Power Cord (4M) Specify -- Belgium Finland, France (250V, 16A)	X	X		X	X						
9906	Power Cord (4M) Specify -- India, Pakistan, S. Africa (250V, 16A)	X	X		X	X						
9907	Power Cord (4M) Specify -- Italy (250V, 10A and 16A)	X	X		X	X						
9904	Power Cord (4M) Specify -- Israel (250V, 6-16A)	X	X		X	X						
9911	Power Cord (4M) Specify -- All (Standard rack power cord)	X	X		X	X						
9004	Southern Hemisphere Designator for Monitors	X	X	X	X	X	X	X	X	X	X	

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3	1 5 0	1 7 0	2 6 5	2 7 5
Rack ID												
4651	Rack Indicator, Rack #1				X	X						
4652	Rack Indicator, Rack #2				X	X						
4653	Rack Indicator, Rack #3				X	X						
4654	Rack Indicator, Rack #4				X	X						
4655	Rack Indicator, Rack #5				X	X						
4656	Rack Indicator, Rack #6				X	X						
4657	Rack Indicator, Rack #7				X	X						
4658	Rack Indicator, Rack #8				X	X						
4659	Rack Indicator, Rack #9				X	X						



Midrange systems

This chapter covers the 7038-6M2 IBM @server pSeries 650 server.

4.1 7038-6M2 IBM @server pSeries 650 server

Figure 4-1 shows the p650, the topic of this section.



Figure 4-1 7038-6M2 p650 with one I/O drawer

The p650 is a mid-range member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM. Positioned between the pSeries 630 and the powerful pSeries 670, the p650 delivers the power, capacity, and expandability required for e-business mission-critical computing.

The p650 is intended for key commercial processing segments such as e-business, ERP, SCM, and Business Intelligence.

For ERP applications in the business processing sector, the p650 is an excellent application server or combined database/application server with its powerful processors, memory capacity, and optional I/O expansion capability.

In e-business environments, the p650 can serve as a fast and highly reliable business-to-business Web server, going to other systems for business data or hosting the data storage itself.

4.1.1 Highlights of the p650

The p650 delivers a cost-efficient growth path to the future with the following characteristics:

- ▶ 64-bit system scalability in 2-, 4-, 6-, and 8-way configurations with the following processor options:
 - 2-way 1.2 GHz POWER4+ with 8 MB shared L3 cache per processor card
 - 2-way 1.45 GHz POWER4+ with 32 MB shared L3 cache per processor card
- ▶ System memory expandable from 2 GB to 64 GB.
- ▶ Rack-mounted drawer utilizes 8U (EIA Units).
- ▶ Customers can expand their capacity with the option of adding up to eight 7311-D10 or 7311-D20 I/O drawers. Supports up to eight 7311-D10 or 7311-D20 I/O drawers per server.
- ▶ Capacity Upgrade on Demand for processors and memory.
- ▶ Logical partitioning (dynamic LPAR with AIX 5L Version 5.2; static LPAR with AIX 5L Version 5.1).
- ▶ Support for SCSI, SSA, and Fibre Channel attached storage systems.
- ▶ Support for 32-bit and 64-bit applications.
- ▶ AIX 5L Version 5.1 or 5.2 license included.

4.1.2 Technical overview of the p650

This section provides a description of the system enhancements and shows a diagram of the p650 system architecture.

System enhancements

The following are recent enhancements to the p650.

RIO-2 Interface

The RIO-2 interface substantially increases the speed of RIO (1 GHz vs 500 MHz), impacting HPC connectivity. The following are the enhancement features:

- ▶ RIO-2 loop adapters - primary (FC 6415) and secondary (FC 6416).
- ▶ RIO-2 enabled system planar.
- ▶ FC 6581 is MES only.
- ▶ FC 9581 is factory only RIO-2 enablement indicator.

- ▶ RIO-2 loop adapter for 7311-D10 I/O drawer (FC 6431).

Split backplane option

The split backplane option creates increased value for partitioning to the second partition. The following are the enhancement features:

- ▶ Optional split backplane and associated cables (FC 6579)
- ▶ Two pairs of hot-swappable disk drive bays
- ▶ Ultra320 SCSI capable
- ▶ Supports disk duplexing and two logical partitions within a p650 server

System architecture

Figure 4-2 shows the diagram of the p650 system architecture.

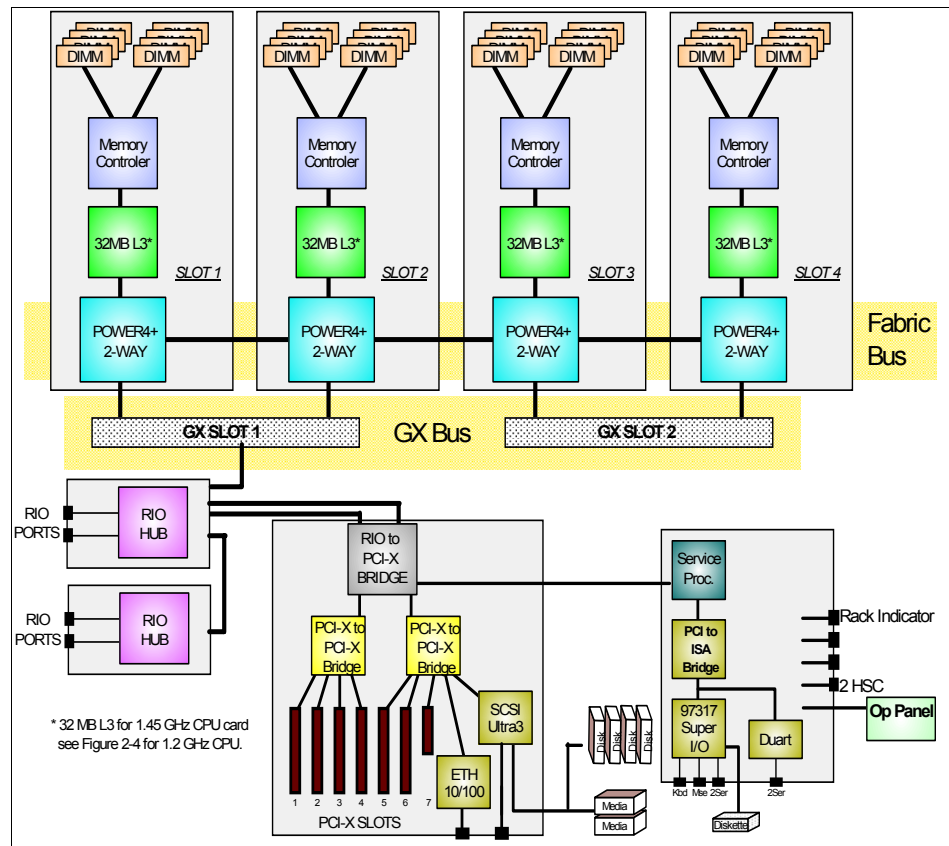


Figure 4-2 p650 system architecture

The p650 processor card contains one Single Chip Module (SCM), L3 cache, and memory.

An SCM contains only one POWER4+ chip (a chip includes two cores) in contrast to Multichip Modules (MCMs), which contain four POWER4+ chips on one module and are used in the pSeries 655, 670, or 690.

The processor chip (consisting of two processors, an L2 cache, and a bus controller), L3 cache, memory, memory controller, and card are all mounted in a rugged metal enclosure collectively named a processor book.

Memory access is through the on-chip L2 cache and L3 cache directory controller to the off-chip L3 cache and finally through the memory controller and synchronous memory interface (SMI) to the memory DIMMs.

There is a fabric bus used for communication between the different SCMs.

The processor I/O path is through the GX bus.

The p650 has three main backplane cards building the internal system buses: GX backplane, the fabric bus backplane, and the I/O subsystem, which is made by a sandwich of two cards (the I/O backplane and the CSP 4 card).

4.1.3 RAS features

The following are the RAS features of the p650.

- ▶ Copper and SOI microprocessors
- ▶ Chipkill ECC and bit-steering memory
- ▶ ECC L2 cache and L3 cache
- ▶ Service processor
- ▶ Hot-swappable disk bays
- ▶ Hot-swappable media bays
- ▶ Hot-plug/blind-swap PCI slots in p650 system unit and 7311-D10 I/O drawer
- ▶ Hot-plug power supplies and cooling fans
- ▶ Dynamic Processor Deallocation
- ▶ Dynamic Processor Sparing
- ▶ Dynamic deallocation of logical partitions and PCI bus slots
- ▶ Redundant power supplies and cooling fans

4.1.4 Minimum and standard features p650

Table 4-1 summarizes the minimum and standard features of the p650.

Table 4-1 p650 minimum and standard features

Specification	Detail
Microprocessor type	64-bit POWER4+.
Min. number of processors	2-way 1.2 GHz POWER4+ with 8 MB shared L3 cache per processor card FC 5122 or 2-way 1.45 GHz POWER4+ with 32 MB shared L3 cache per processor card FC 5208.
Memory (minimum)	2 GB.
Data/Instruction (L1) cache	32 KB/64 KB per processor core.
Level 2 (L2) cache	1.5 MB per processor card on multichip module.
Level 3 (L3) cache	<ul style="list-style-type: none"> ▶ 32 MB L3 cache on 1.45 GHz processor card FC 5208 ▶ 8 MB L3 cache on 1.2 GHz processor card FC 5122
Internal disk bays	Four hot-swappable disk drive bays or optional two independent pairs of disk drives are available using a split backplane.
Internal disk drives	Each bay can contain one of the following disks <ul style="list-style-type: none"> ▶ 36.4 GB Ultra320 10K RPM (FC 3273) ▶ 36.4 GB Ultra320 15K RPM (FC 3277) ▶ 73.4 GB Ultra320 10K RPM (FC 3274) ▶ 73.4 GB Ultra320 15K RPM (FC 3278) ▶ 146.8 GB Ultra320 10K RPM (FC 3275)
CD-ROM, DVD-RAM, or DVD-ROM drive	At least one CD-ROM or DVD-RAM must be configured on an initial order: <ul style="list-style-type: none"> ▶ 16/48X DVD-ROM auto-docking module (FC 2635) ▶ 40X CD-ROM auto-docking module (FC 2628) ▶ 4.7 GB R/W DVD-RAM auto-docking module (FC 2629)
Media bays	Two (one available) hot-swappable.

Specification	Detail
PCI slots	Seven blind-swap hot-pluggable PCI slots: <ul style="list-style-type: none"> ▶ Six 64-bit 133 MHz 3.3 volt PCI-X slots, full length, blind-swap hot-plug and ▶ One 32-bit 66 MHz 3.3 volt PCI-X slot, half length, blind-swap hot-plug
Standard Ports	<ul style="list-style-type: none"> ▶ Four serial ports ▶ Keyboard and mouse ports ▶ Two Hardware Management Console ports ▶ One Rack Status Beacon port
Integrated SCSI adapters	Two integrated SCSI controllers (Ultra3).
Integrated Ethernet adapter	One Ethernet 10/100 controller.
Diskette drive	Yes.
System support (service) processor	Yes.
AIX operating system version	AIX 5L Version 5.1 or 5.2 license included.
Capacity Upgrade on Demand	Offers flexibility to cost effectively and easily add permanent processing and memory capacity to help meet workload growth.
On/Off Capacity on Demand	Provides temporary processor use to meet unexpected workload demands.
Logical Partitioning (LPAR)	Offers greater flexibility in using available capacity and dynamically matching resources to changing business needs. A fully configured pSeries 650 can support up to eight partitions. Refer to 16.1, "Partitioning" on page 686 for more information.
Warranty	One year.

4.1.5 System expansion of the p650

Table 4-2 shows the possible maximum processor, memory, and storage expansion features.

Table 4-2 Possible maximum processor, memory, and storage features

System expansion	
SMP configuration	From a 2-way (1.2 GHz FC 5122 or 1.45 GHz FC 5208) to an 8-way. In 2-, 4-, 6-, and 8-way SMP configurations (one, two or three additional processor cards).
L3 cache per processor card	<ul style="list-style-type: none"> ▶ 32 MB L3 cache on 1.45 GHz processor card FC 5208 ▶ 8 MB L3 cache on 1.2 GHz processor card FC 5122
Memory	<p>Up to 64 GB using the following available order combinations:</p> <ul style="list-style-type: none"> ▶ FC 4452 2048 MB (4 x 512 MB) 208-pin 8 ns DDR SDRAM DIMMs ▶ FC 4453 4096 MB (4 x 1024 MB) 208-pin 8 ns stacked DDR SDRAM DIMMs ▶ FC 4454 8192 MB (4 x 2048 MB) 208-pin 8 ns stacked DDR SDRAM DIMMs
PCI expansion slots	<p>Up to 48 additional hot-plug/blind-swap adapters (40 64-bit PCI-X; eight 64-bit PCI) using up to eight 7311-D10 I/O drawers (six adapters each).</p> <p>Up to 56 additional hot-plug adapters (64-bit PCI-X) using up to eight 7311-D20 I/O drawers (seven adapters each).</p>
Disk drive and disk bay expansion	<p>Four hot-swappable disk drive bays with 36.4 GB to 587.2 GB or optional two independent pairs of disk drives with 36.4 GB to 293.6 GB of internal disk storage are available using a split backplane.</p> <p>Up to 96 hot-swappable via up to eight 7311-D20 I/O drawers (12 bays each); Up to 14.0 TB of additional disk storage available.</p>

4.1.6 Features of the p650

This section lists the internal features that can be added to the p650 configuration. The status of a feature is indicative of these qualifications:

- X** Available. Indicates features that are available and are orderable on the specified model.
- R** Must be removed. Indicates this feature is not supported and must be removed during a model conversion.
- W** Withdrawn. Indicates a feature that is no longer available to order.

Features not listed in the provided categories, indicate that the feature is not supported on this model. Some categories, such as keyboards, mice, language specify codes and power cords are not included.

Table 4-3 shows the features of p650.

Table 4-3 Features of the p650

Feature code	Description	7038 6M2
Host attachment adapters		
8398	SP Switch2 PCI-X Attachment Adapter	X
Asynchronous adapters		
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI bus	X
2944	128-Port Asynchronous Controller, PCI bus	X
Asynchronous cables		
2936	Asynchronous Cable EIA-232/V.24	X
2934	Asynchronous Terminal/Printer Cable EIA-232	X
3926	Asynch Printer/Terminal Cable, 9-pin to 25-pin, 4M	X
8131	128-Port Asynchronous Controller Cable, 4.5 Meter	X
8132	128-Port Asynchronous Controller Cable, 23cm (9-in.)	X
8121	Attachment Cable, Hardware Management Console to Host, 15-Meter	X
8120	Attachment Cable, Hardware Management Console to Host, 6-Meter	X
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232	X

Feature code	Description	7038 6M2
8133	RJ-45 to DB-25 Converter Cable	X
8136	Rack Mountable Remote Asynchronous Node 16-Port EIA-232	X
3925	Serial Port Converter Cable, 9-Pin to 25-Pin	X
SCSI adapters		
6204	PCI Universal Differential Ultra SCSI Adapter	X
6203	PCI Dual Channel Ultra3 SCSI Adapter	X
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter	X
5711	PCI-X Dual Channel Ultra320 SCSI RAID Adapter (in cassette)	X
5710	PCI-X Dual Channel Ultra320 SCSI Adapter (in cassette)	X
Serial adapters		
6231	128 MB DRAM Option Card	X
6230	Advanced SerialRAID Plus Adapter	X
6235	32 MB Fast-Write Cache Option Card	X
Fibre Channel adapters		
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus	W
6239	2 Gigabit Fibre Channel PCI-X Adapter	X
Fiber cables		
2456	LC-SC 50 Micron Fiber Converter Cable	X
2459	LC-SC 62.5 Micron Fiber Converter Cable	X
Graphics accelerators		
2848	POWER GXT135P Graphics Accelerator	X
2849	POWER GXT135P Graphics Accelerator with Digital Support	X
Display cable		
4242	6 Foot Extender Cable for Displays (15 pin D-shell to 15 pin D-shell)	X

Feature code	Description	7038 6M2
LAN adapters		
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	X
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	X
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	X
4959	IBM Token-Ring PCI Adapter	X
2946	Turboways 622 Mbps PCI MMF ATM Adapter	X
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	X
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	X
4962	10/100 Mbps Ethernet PCI Adapter II	X
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	X
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	X
WAN adapters		
2962	2-Port Multiprotocol PCI Adapter	X
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	X
6313	Quad Digital Trunk Telephony PCI Blind Swap Adapter	X
WAN cables		
2861	ARTIC960Hx 4-Port EIA-232 Cable	X
2864	ARTIC960Hx 4-Port V.35 (DTE) Cable	X
2863	ARTIC960Hx 4-Port X.21 Cable	X
2710	ARTIC960Hx/RxD 4-Port E1 RJ-45 Cable	X
2709	ARTIC960Hx/RxD 4-Port T1 RJ-45 Cable	X
2951	Cable, V.24/EIA-232	X
2952	Cable, V.35	X
2953	Cable, V.36/EIA-499	X
2954	Cable, X.21	X

Feature code	Description	7038 6M2
Adapter mounting hardware		
4599	PCI Blind Swap Cassette Kit, Single Wide Adapters, Universal	X
4597	PCI Blind Swap Cassette Kit, Single Wide Adapters, Short	X
Miscellaneous		
4963	PCI Cryptographic Coprocessor (FIPS-4)	X
4960	IBM e-business Cryptographic Accelerator	X
2737	Keyboard/Mouse Attachment Card - PCI	X
2877	IBM ARTIC960RxD Quad DTA, H.100, 4-Drop Cable	X
4353	H-100 Bus 8-position Cable	X
4692	Rack Status Beacon Cable, Junction Box Daisy Chain	X
4691	Rack Status Beacon Cable, Junction Box To Drawer Or Status Beacon	X
3124	Serial to Serial Port Cable for Drawer/Drawer	X
3125	Serial to Serial Port Cable for Rack/Rack	X
9004	Southern Hemisphere Designator for Monitors	X
Miscellaneous internal system		
6412	Remote I/O Loop Adapter, Secondary	X
6411	Remote I/O Loop Adapter, Primary	X
6416	RIO-2 Remote I/O Loop Adapter, Secondary	X
6415	RIO-2 Remote I/O Loop Adapter, Primary	X
SCSI internal		
4262	Internal Cable Assembly For Split Backplane	X
4255	Internal Cable, Integrated Controller To SCSI Backplane	X
4264	Internal Cable, Integrated Controller To Media Backplane	X
4261	Internal Cable, Media Backplane To Ultra3 SCSI Backplane (4-Pack)	X

Feature code	Description	7038 6M2
SCSI external		
2424	0.6M 16-bit SCSI-2 System to System Cable	X
2425	2.5M 16-bit SCSI-2 System to System Cable	X
4265	External Cable For Split Backplane	X
2118	Converter Cable, VHDCI to P, Mini-68 pin to 68 pin, 0.3M	X
2114	PCI SCSI Adapter 16-Bit Differential External Y Cable	X
Rack related		
6001	Power Control Cable (SPCN) - 2 meter	X
6006	Power Control Cable (SPCN) - 3 meter	X
6008	Power Control Cable (SPCN) - 6 meter	X
6007	Power Control Cable (SPCN) - 15 meter	X
3146	RIO-2 (Remote I/O-2) Cable, 1.2M	X
3147	RIO-2 (Remote I/O-2) Cable, 3.5M	X
3148	RIO-2 (Remote I/O-2) Cable, 10M	X
SCSI disks		
3158	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W
3273	36.4 GB 10K RPM Ultra320 SCSI Disk Drive Assembly	X
3277	36.4 GB 15K RPM Ultra320 SCSI Disk Drive Assembly	X
3159	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W
3274	73.4 GB 10K RPM Ultra320 SCSI Disk Drive Assembly	X
3278	73.4 GB 15K RPM Ultra320 SCSI Disk Drive Assembly	X
3275	146.8 GB 10K RPM Ultra320 SCSI Disk Drive Assembly	X
Monitors		
3636	L200p Flat Panel Monitor	X
3635	T210 Flat-Panel Monitor	W
3628	IBM P260/P275 Color Monitor, Business Black, and Cable	X

Feature code	Description	7038 6M2
3627	IBM P76/P77 Color Monitor, Business Black, Captured Cable	W
3637	IBM T541H 15" TFT Color Monitor, Business Black, Captured Cable	X
DIMMS		
4452	2048 MB (4x512 MB) DIMMs, 208-pin, 8NS DDR SDRAM	X
4453	4096 MB (4x1024 MB) DIMMs, 208-pin, 8NS Stacked DDR SDRAM	X
4454	8192 MB (4x2048 MB) DIMMs, 208-pin, 8NS Stacked DDR SDRAM	X
8052	4096 MB (4x1024 MB) DIMMs, Express Configuration, Factory Only	X
Memory CUoD		
7052	4096 MB CUoD Memory Activation For Feature #7057	X
7053	4096 MB CUoD Memory Activation For Feature #7058	X
7057	4096 MB (4x1024 MB) DIMMs, CUoD, 0 MB Active	X
7058	8192 MB (4x2048 MB) DIMMs, CUoD, 0 MB Active	X
Drawer related		
6287	AC Power Supply, 1100 W	X
9172	Power Specify, AC	X
SCSI CD/DVD		
2635	16X/48X(max) SCSI Auto-docking DVD-ROM Drive	X
2628	48X (Max) SCSI-2 Internal Auto-docking CD-ROM Drive	X
2629	4.7 GB SCSI-2 Auto-docking DVD-RAM Drive	X
SCSI tapes		
6185	20/40 GB 4mm Internal Auto-docking Tape Drive	X
6131	60/150 GB 16-bit 8mm Internal Auto-docking Tape Drive	X
6169	IBM 80/160 GB Internal Auto-docking Tape Drive with VXA Technology	X

Feature code	Description	7038 6M2
Media mounting hardware		
4251	Backplane For Auto-Docking Media Drives	X
9580	CUoD Smart Card Indicator, ISU Only	X
6581	RIO-2 Enabled System Planar	X
9581	RIO-2 Enablement Indicator, Factory Only	X
8841	Mouse - Business Black with Keyboard Attachment Cable	X
Bundle/pricing features		
8039	Entitlement Feature For Express Configuration -- pSeries 650 (201M)	X
8048	Entitlement Feature For Express Configuration -- pSeries 650 (401M)	X
8117	Entitlement Feature For Express Configuration -- pSeries 650 (451M)	X
8118	Entitlement Feature For Express Configuration -- pSeries 650 (651M)	X
8049	Entitlement Feature For Express Configuration -- pSeries 650 (801M)	X
8119	Entitlement Feature For Express Configuration -- pSeries 650 (851M)	X
8053	Entitlement Feature For Express Configuration -- pSeries 650 (200M)	X
8054	Entitlement Feature For Express Configuration -- pSeries 650 (400M)	X
8057	Entitlement Feature For Express Configuration -- pSeries 650 (450M)	X
8058	Entitlement Feature For Express Configuration -- pSeries 650 (650M)	X
8055	Entitlement Feature For Express Configuration -- pSeries 650 (800M)	X
8059	Entitlement Feature For Express Configuration -- pSeries 650 (850M)	X

Feature code	Description	7038 6M2
Disk bays		
6579	Split SCSI Backplane For Hot-swap Disks (Split 4-pack)	X
6578	Ultra3 SCSI Backplane for Hot-swap Disks (4-pack)	X
Software preload		
5005	Software Preinstall	X
Routing indicators		
0986	CCS Customer Service Specify (US)	X
5001	Customer Service Specify	X
Processors		
5122	2-Way 1.2 GHz POWER4+ Processor Card	X
5208	2-Way 1.45 GHz POWER4+ Processor Card	X
8050	2-way 1.2 GHz POWER4+ Processor Card, Express Configuration, Factory Only	X
8051	2-way 1.45 GHz POWER4+ Processor Card, Express Configuration, Factory Only	X
Processor CUoD		
7011	Two Processor Activation for CUoD Processor Feature 7014	X
7013	On/Off CoD 30-Day Two Processor Activation for Processor Feature #7014	X
7014	2-Way 1.45 GHz POWER4 CUoD Processor Card, 0-Way Active	X
Processor fabric		
5120	Processor Card Backplane For 4-Way Configuration	X
5123	Processor Card Backplane For 6-Way Configuration	X
5124	Processor Card Backplane For 8-Way Configuration	X
4693	Rack Status Beacon Junction Box	X

Feature code	Description	7038 6M2
Rack ID		
4651	Rack Indicator, Rack #1	X
4652	Rack Indicator, Rack #2	X
4653	Rack Indicator, Rack #3	X
4654	Rack Indicator, Rack #4	X
4655	Rack Indicator, Rack #5	X
4656	Rack Indicator, Rack #6	X
4657	Rack Indicator, Rack #7	X
4658	Rack Indicator, Rack #8	X
4659	Rack Indicator, Rack #9	X

4.1.7 Configuration notes

The following section covers the configuration notes for the p650.

Processors

The following are the configuration notes regarding the processors:

- ▶ The p650 system contains four processor slots. The available 2-way processor cards can be combined to create 2-way, 4-way, 6-way, or 8-way systems.
- ▶ When two or more processor cards are installed, all cards must have the same processor frequency.
- ▶ There are eight memory DIMM positions on each processor card.

Memory

The following are the configuration notes regarding the memory:

- ▶ Memory is installed in groups of four DIMMs mounted on the processor cards. Each processor card has 8 DIMM slots.
- ▶ Each processor card must have an equal amount of memory installed. Balancing memory across the installed processor cards allows memory accesses in a coordinated parallel manner and can be utilized to provide optimal performance.

- ▶ A system with a single processor card (2-way) may have a maximum of 16 GB of memory based on the maximum memory feature available (FC 4454, 4 x 2048 MB). The maximum memory on a 4-way is 32 GB, on a 6-way is 48 GB, and on an 8-way is 64 GB.
- ▶ Initial orders require identical combinations of memory on each processor card. Active and CUoD memory features are counted as equal. MES orders may use any combination of supported memory features as long as the total memory capacity (active plus CUoD) on each processor card is the same.
- ▶ Plans for future memory upgrades should be taken into account to obtain optimal performance without limiting future memory expansion.
- ▶ IBM recommends systems configured with 4 GB of memory or greater to have access to a 4-mm or 8-mm tape drive for submission of system dump information, if required. This function may be accomplished using a locally-attached or network-attached devices as appropriate.

Table 4-4 shows the recommended memory configurations.

Table 4-4 Recommended memory configurations

Total mem.	One proc. card	Two processor cards		Three processor cards			Four processor cards			
		1st	2nd	1st	2nd	3rd	1st	2nd	3rd	4th
2 GB	2									
4 GB	4	2	2							
6 GB	4 + 2			2	2	2				
8 GB	4 + 4	4	4				2	2	2	2
10 GB	8 + 2									
12 GB	8 + 4	4 + 2	4 + 2	4	4	4				
16 GB	8 + 8	4 + 4	4 + 4				4	4	4	4
18 GB				4 + 2	4 + 2	4 + 2				
20 GB		8 + 2	8 + 2							
24 GB		8 + 4	8 + 4	8	8	8	4 + 2	4 + 2	4 + 2	4 + 2
30 GB				8 + 2	8 + 2	8 + 2				
32 GB		8 + 8	8 + 8				8	8	8	8
36 GB				8 + 4	8 + 4	8 + 4				

Total mem.	One proc. card	Two processor cards		Three processor cards			Four processor cards			
		1st	2nd	1st	2nd	3rd	1st	2nd	3rd	4th
40 GB							8 + 2	8 + 2	8 + 2	8 + 2
48 GB				8 + 8	8 + 8	8 + 8	8 + 4	8 + 4	8 + 4	8 + 4
64 GB							8 + 8	8 + 8	8 + 8	8 + 8

Power

Two power supplies must be ordered to provide redundancy for enhanced system availability.

Power cords

For information, refer to Appendix C, “Power cord features” on page 907.

Racks

The following are the configuration notes regarding the racks:

- ▶ The p650 systems CEC and I/O drawers are packaged as rack-mounted devices and are supported in the 7014-T00 and T42 racks. An existing 7014 rack can be used for the p650 if sufficient space and power is available.

Important: AC-powered 6M2 systems require power distribution features type 7: FC 7176, FC 7177, FC 7178, FC 9176, FC 9177, or FC 9178.

The AC power requirements for the p650 server are significantly higher than previous IBM rack mounted servers and require use of the new type 7 PDU. The Type7 PDU has 9 IEC320-C13 outlets configured in groups of three.

- ▶ An IBM 7014-T00 or T42 rack must have at least one Power Distribution Unit (PDU) per pSeries 650 system installed. It is recommended that each power supply in a pSeries 650 system be connected to a different PDU. No more than two pSeries 650 power supplies may be connected to the same PDU. Each PDU has nine C13 power connectors, in groups of three. When a pSeries 650 power supply is connected to one PDU power outlet, the other two connectors in that group of three may not be used.
- ▶ The maximum quantities of 7038-6M2 systems supported in 7014 racks are as follows:
 - 7014-T00: The maximum is four.

- 7014-T42: The maximum is five (requires a fifth PDU mounted horizontally in the rack).
- ▶ Supported devices, such as SCSI or SSA disk subsystems or tape subsystems, may share the rack with the p650 system.

I/O drawers

The following are the configuration notes regarding the I/O drawers:

- ▶ p650 systems can be attached to up to eight 7311-D10 or 7311-D20 I/O drawers.
- ▶ Each 7311-D10 I/O drawer provides six PCI slots. Refer to the 7311-D10. Refer to 5.2, “7311-D10 I/O Drawer” on page 206 for specifications.
- ▶ Each 7311-D20 I/O drawer provides seven PCI slots and up to twelve disk drives. Refer to 5.1, “7311-D20 Expansion Drawer” on page 200 for specifications.

I/O drawer attachment

The following are the configuration notes regarding the I/O drawer attachment:

- ▶ The p650 system allows up to eight I/O drawers to be attached to the system CEC drawer. I/O drawers should 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 p650 drawer via the following cables:
 - RIO cables for data transfer
 - 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. p650 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.
- ▶ RIO ports are provided by remote I/O loop adapter cards. The first primary RIO adapter (FC 6411 or FC 6415) is required. An optional secondary RIO adapter (FC 6412 or FC 6416) provides two additional RIO ports. The first secondary RIO adapter must be installed prior to installing the second primary RIO adapter.
- ▶ Up to four RIO loops are available, each supporting up to four I/O drawers and up to a maximum of eight 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 in

various lengths to attach I/O drawers within a single rack or across multiple racks, if desired.

- FC 3146 = 1.2 meters
- FC 3147 = 3.5 meters
- FC 3148 = 10 meters
- ▶ The configurator validates system connections when the RIO cables are ordered on the individual 7311-D10 or 7311-D20 drawers. If RIO cables are ordered on the 7038-6M2 server, system validation cannot be performed.
- ▶ Power control for the I/O drawers is provided using one loop. The number of power control (SPCN) cables required is equal to one plus the number of I/O drawers attached to the system. A minimum of two power control cables (FC 6001 = 2 meters, FC 6006 = 3 meters, FC 6007 = 15 meters or FC 6008 = 6 meters) are required for attachment of the first drawer. Each additional drawer requires one additional power control cable to complete the loop attachment.
- ▶ The number of RIO or RIO-2 adapters must be less than or equal to the number of 2-way processor cards.
- ▶ Each RIO or RIO-2 adapter supports one RIO loop. Each RIO loop supports up to two 7311-D20 I/O drawers or up to four 7311-D10 I/O drawers.
- ▶ 7311-D10 and 7311-D20 I/O drawers are not supported in the same I/O loop.
- ▶ For systems with both 7311-D10 and 7311-D20 I/O drawers attached, the combined maximum number of I/O drawers is eight.
- ▶ To achieve the higher remote I/O performance of RIO-2, the system must have a RIO-2 Enabled System Planar (FC 9581 or FC 6581). Also, to achieve RIO-2 performance, all adapters in a loop must be RIO-2. RIO and RIO-2 adapter combinations are supported in the same loop; however, the loop will operate at the standard RIO speed.

Disks, media, and boot devices

The following are the configuration notes regarding the disk, media, and boot devices:

- ▶ Boot is supported using SCSI adapters, SSA adapters, Fibre Channel adapters, or from network using Ethernet or token-ring adapters.
- ▶ A CD-ROM, DVD-ROM, or DVD-RAM is required for each system. One bay is available for an additional media device.
- ▶ The minimum configuration requires at least one disk drive having a capacity of 36.4 GB or greater.

- ▶ Systems ordered with a single 4-pack SCSI backplane require these features for proper cabling and operation: FC 6578, FC 4255, FC 4251, and FC 4261. The SCSI disk backplane (FC 6578) must be connected to the integrated internal SCSI port with cable FC 4255. The FC 4251 media backplane must be connected to the FC 6578 disk backplane with cable FC 4261.

PCI slots and adapters

The following are the configuration notes regarding the PCI slots and adapters:

- ▶ The p650 has seven PCI-X slots. All slots operate at 3.3 volts. There are six full-length 64-bit slots and one short 32-bit slot. Refer to Appendix B, “Adapter placement guidelines” on page 785 for more information.
- ▶ The p650 PCI slot population rules are complex. Extensive configuration rules and checking procedures are incorporated into the ECFGRS6000 configurator to ensure valid system configurations. Configurations generated without utilizing the ECFGRS6000 configurator may create orders that cannot be manufactured, resulting in possible order rejection and delayed delivery.
- ▶ System maximum limits for adapters and devices may not provide optimal system performance. These limits are given for connectivity and functionality assurance.
- ▶ When the p650 is attached to one or more I/O drawers, additional rules apply for some PCI adapters regarding the maximum supported adapter quantities for the combined system. For adapters not shown in the list below, the combined maximum is simply the number supported in a 6M2 server, plus eight times the number supported in a D10 I/O drawer. For further information regarding adapter placement rules, refer to Appendix B Adapter Placement Guidelines”

The adapters in Table 4-5 are limited to the combined maximum quantity shown.

Table 4-5 Max. number of adapters supported in combined system 6M2

Feature	Description	Quantity
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI bus	32
2944	28-Port Asynchronous Controller, PCI bus	32
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter	30
6203	PCI Dual Channel Ultra3 SCSI Adapter	30
6204	PCI Universal Differential Ultra SCSI Adapter	40
6230	Advanced SerialRAID Plus Adapter	32
6231	128 MB DRAM Option Card	32

Feature	Description	Quantity
6235	32 MB Fast-Write Cache Option Card	32
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus	32
2848	POWER GXT135P Graphics Accelerator	8
2946	Turboways 622 Mbps PCI MMF ATM Adapter	30
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	30
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	30
4959	IBM Token-Ring PCI Adapter	30
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	20
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	30
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	30
2737	Keyboard/Mouse Attachment Card - PCI	8
4960	IBM e-business Cryptographic Accelerator	15
4963	PCI Cryptographic Coprocessor (FIPS-4)	32
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	30
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	30
8398	SP Switch2 PCI-X Attachment Adapter	2

The following adapters in Table 4-6 cannot be installed in slot 7 of the p650. The maximum number of any combination of these adapters is six.

Table 4-6 Adapters that cannot be installed in slot 7

Feature	Description
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter
6203	PCI Dual Channel Ultra3 SCSI Adapter
6230	Advanced SerialRAID Plus Adapter
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus
2946	Turboways 622 Mbps PCI MMF ATM Adapter
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter

Feature	Description
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter
4961	IBM Universal 4-Port 10/100 Ethernet Adapter
5700	IBM Gigabit Ethernet-SX PCI-X Adapter
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter
8398	SP Switch2 PCI-X Attachment Adapter
6313	Quad Digital Trunk Telephony PCI Blind Swap Adapter

Hot-plug options

The following are the configuration notes regarding the Hot-Plug options:

- ▶ It is not necessary to power down the system to install/remove/replace certain hot-plug options. The following options are hot-plug capable:
 - System power supplies (assumes redundant power supplies installed)
 - Most PCI adapters
 - Information regarding hot-plugging procedures is contained in the *pSeries 650 Installation Guide and Service Guide*, found at:
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm
 - Any PCI adapter supporting the system boot device or system console is not hot-plugged.
- ▶ The following adapters are not hot-plug capable:
 - POWER GXT135P Graphics Accelerator (FC 2848)
 - 2-Port Multiprotocol PCI Adapter (FC 2962)

Logical partitioning (LPAR)

Refer to 16.1, “Partitioning” on page 686.

Physical specifications

Refer to Appendix A, “Site and hardware planning information” on page 747.

4.1.8 Express configurations

The following are the express configurations available for the p650. Additional, optional 7038-6M2 features may be added as desired. Except as noted, all standard configuration rules for p650 remain unchanged. Table 4-7 provides the express configuration summary for the p650.

Table 4-7 p650 express configuration summary

Express config.	OS	CPU	Memory	DASD
200 M	AIX 5L	One 2-way 1.2 GHz POWER4+	4096 MB	Two 146.8 GB
201 M	Linux	One 2-way 1.2 GHz POWER4+	4096 MB	Two 146.8 GB
400 M	AIX 5L	Two 2-way 1.2 GHz POWER4+	8192 MB	Two 146.8 GB
401 M	Linux	Two 2-way 1.2 GHz POWER4+	8192 MB	Two 146.8 GB
800 M	AIX 5L	Four 2-way 1.2 GHz POWER4+	16384 MB	Two 146.8 GB
801 M	Linux	Four 2-way 1.2 GHz POWER4+	16384 MB	Two 146.8 GB
450 M	AIX 5L	Two 2-way 1.45 GHz POWER4+	8192 MB	Two 146.8 GB
451 M	Linux	Two 2-way 1.45 GHz POWER4+	8192 MB	Two 146.8 GB
650 M	AIX 5L	Three 2-way 1.45 GHz POWER4+	12288 MB	Two 146.8 GB
651 M	Linux	Three 2-way 1.45 GHz POWER4+	12288 MB	Two 146.8 GB
850 M	AIX 5L	Four 2-way 1.45 GHz POWER4+	16384 MB	Two 146.8 GB
851 M	Linux	Four 2-way 1.45 GHz POWER4+	16384 MB	Two 146.8 GB

Express Configuration - pSeries 650 (200M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (200M) (FC 8053)
- ▶ One 2-way, 1.2 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8050)
- ▶ One 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each with 146.8 GB capacity (FC 3275)

Express Configuration - pSeries 650 (400M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (400M) (FC 8054)
- ▶ Two 2-way, 1.2 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8050)
- ▶ Two 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each with 146.8 GB capacity (FC 3275)

Express Configuration - pSeries 650 (800M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (800M) (FC 8055)
- ▶ Four 2-way, 1.2 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8050)
- ▶ Four 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each with 146.8 GB capacity (FC 3275)

Express Configuration - pSeries 650 (450M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (450M) (FC 8057)
- ▶ Two 2-way, 1.45 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8051)

- ▶ Two 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each with 146.8 GB capacity (FC 3275)

Express Configuration - pSeries 650 (650M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (650M) (FC 8058)
- ▶ Three 2-way, 1.45 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8051)
- ▶ Three 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each with 146.8 GB capacity (FC 3275)

Express Configuration - pSeries 650 (850M)

One 7038-6M2 system, including a minimum of:

- ▶ One Entitlement Feature For Express Configuration - pSeries 650 (850M) (FC 8059)
- ▶ Four 2-way, 1.45 GHz POWER4+ Processor Card, Express Configuration, Factory Only (FC 8051)
- ▶ Four 4096 MB (4 x 1024 MB) DIMMs, Express Configuration, Factory Only (FC 8052)
- ▶ Two disk drives, each having 146.8 GB capacity (FC 3275)

The following are Linux ready Express Configurations equivalent of the AIX based Express Configurations:

- ▶ (FC 8039) - Entitlement Feature For Express Configuration - pSeries 650 (201M)
- ▶ (FC 8048) - Entitlement Feature For Express Configuration - pSeries 650 (401M)
- ▶ (FC 8117) - Entitlement Feature For Express Configuration - pSeries 650 (451M)
- ▶ (FC 8118) - Entitlement Feature For Express Configuration - pSeries 650 (651M)
- ▶ (FC 8049) - Entitlement Feature For Express Configuration - pSeries 650 (801M)
- ▶ (FC 8119) - Entitlement Feature For Express Configuration - pSeries 650 (851M)



IBM rack-mounted expansion drawers

This chapter describes the I/O expansion drawers available on IBM @server pSeries 630 Model 6C4 and IBM @server pSeries 650.

The I/O expansion drawers are the following:

- ▶ 7311-D20 IBM Rack-Mounted High-Density Expansion Drawer
- ▶ 7311-D10 IBM I/O Drawer Rack-Mountable Expansion Cabinet

5.1 7311-D20 Expansion Drawer

This section discusses the 7311-D20 I/O drawer.

Figure 5-1 shows the 7311-D20 I/O drawer.



Figure 5-1 7311-D20 I/O drawer

5.1.1 7311-D20 I/O drawer overview

The IBM 7311 Model D20 is a rack-mounted, high density expansion drawer that attaches to the pSeries 630 Model 6C4 and pSeries 650 to provide remote I/O. There are seven hot-swap PCI-X I/O slots in the 7311-D20 and twelve optional hot-swap disk drive bays.

The 7311-D20 occupies 4U (7.00-in.) of space in a rack and mounts in a 19-in. standard rack drawer. The Model D20 is 24 inches deep. The fans, power supplies, and PCI adapters, are top-accessible, while the disk drives are front-accessible for easy service and maintenance.

Figure 5-2 shows the front and rear views of the 7311-D20 I/O drawer.




Figure 5-2 7311 Model D20 front and rear view

The hot-plug mechanism is the same as on the Models 6F1, 6H1, or 6M1; therefore, PCI cards are inserted from the top of the I/O drawer down into the slot. The installed adapters are protected by plastic separators, designed to prevent grounding and damage when adding or removing adapters. Figure 5-3 shows an inside view of the 7311 Model D20 I/O drawer.




Figure 5-3 Inside view of 7311 Model D20 I/O drawer

Both the Model 6C4 and p650 provide RIO or RIO-2 ports for I/O drawer connection either in single or multiple loop (p650 only). Each RIO or RIO-2 port operates at 500 MHz (RIO) or 1 GHz (RIO-2) in bidirectional mode and is capable of passing up to eight bits of data in each direction on each cycle of the RIO or RIO-2 port. If only one I/O drawer is connected, the RIO loop runs in double barrel mode, meaning both cables are used by this I/O drawer, giving a total bandwidth of 2 GBps in RIO mode or 4 GBps in RIO-2 mode for this drawer. One RIO or RIO-2 port (capable of 1 GBps or 2 GBps) will be assigned to a PHB (slots 1-4) and the other RIO or RIO-2 port (capable of 1 GBps or 2 GBps) will be assigned to the other PHB (slots 5-7). If two I/O drawers are connected, in normal operation each I/O drawer uses one cable, giving a total bandwidth of 1 GBps in RIO mode or 2 GBps in RIO-2 mode for each I/O drawer to be shared between the internal devices in that drawer.

The p630 to achieve the higher Remote I/O performance of RIO-2 must have a RIO-2 Enabled System Planar (FC 9581) and 7311-D20 must have RIO-2 riser adapter (FC 6417).

The PCI-X host bridge inside the I/O drawer provides two primary 64-bit PCI-X buses running at 133 MHz. Therefore, a maximum bandwidth of 1 GBps is provided by each of the buses.

Each primary PCI-X bus is connected to a PCI-X-to-PCI-X bridge. The first PCI-X-to-PCI-X bridge provides four slots and the second PCI-X-to-PCI-X bridge provides three slots, both with Extended Error Handling (EEH) for error recovering. All slots are PCI-X slots that operate at 133 MHz and 3.3 volt signaling.

Figure 5-4 shows a conceptual diagram of the 7311-D20 I/O drawer.

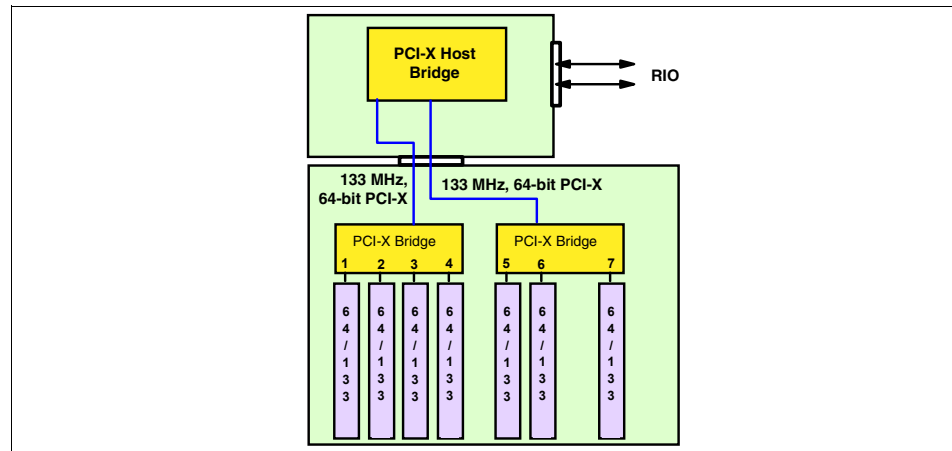


Figure 5-4 Conceptual diagram of the I/O drawer

For power control and monitoring of the I/O drawers, SPCN cables are used. The SPCN cables form a loop similar to the way the RIO cables do. The cabling starts from SPCN port 0 on the CEC to SPCN port 0 on the first I/O drawer connecting them from port 1 to port 0 of second I/O drawer or back to the CEC.

The RIO cabling works in a similar way. The CEC connects from RIO port 0 to port 0 on the first I/O drawer. From port 1 on the I/O drawer, a second I/O drawer could be connected on port 0. The last I/O drawer in the RIO loop connects from port 1 back to port 1 at the CEC.

Figure 5-5 on page 203 shows the cabling of the RIO cables as well as the cabling of SPCN.

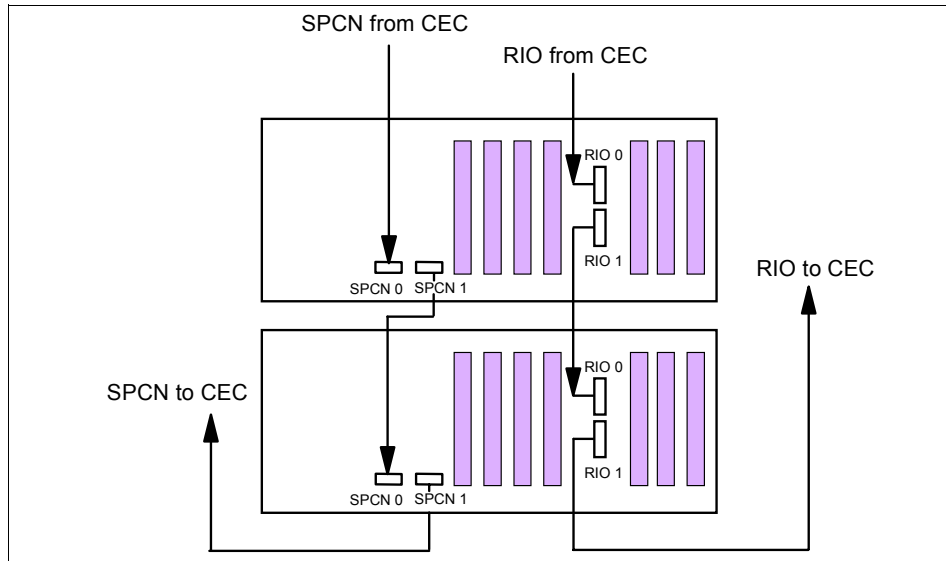


Figure 5-5 Remote I/O and SPCN cabling

As an option, you can add up to two 7311-D20 I/O drawers in the p630 Model 6C4 and eight 7311-D20 I/O drawers in the p650. Drawers can be intermixed on a single server, provided each RIO loop contains a single drawer type.

5.1.2 Highlights

The 7311-D20 has the following characteristics:

- ▶ Two RIO or RIO-2 (2x speed RIO) ports, redundant hot-plug power, redundant hot-plug cooling, and two SPCN ports.
- ▶ Seven PCI-X slots, 3.3 volt keyed, 133 MHz blind-swap hot-plug
- ▶ Up to twelve hot-swappable disks.

Table 5-1 summarizes the features of the 7311-D20.

Table 5-1 7311-D20 I/O drawer at a glance

Features	Detail
Disk drive bays	12 optional hot-swap disk drive bays consisting of two 6-packs
Internal disk drives	18.2, 36.4, 73.4, and 146.8 GB Ultra3 SCSI Disk Drives
PCI slots	7 hot-swap PCI-X 64-bit, 133MHz, 3.3 volt I/O slots

Features	Detail
Hardware minimum requirements	<ul style="list-style-type: none"> ▶ I/O Riser card: <ul style="list-style-type: none"> – (FC 6413) - RIO Ports to I/O Planar Riser Card or – (FC 6417) - RIO-2 Ports to I/O Planar Riser Card ▶ Power supply (FC 6268) ▶ Power specify (FC 9172) ▶ Power cord specify (FC 98xx, FC 99xx) ▶ Language group specify (FC 9300 or FC 97xx) ▶ Two Remote I/O cables (FC 3147 or FC 3148) ▶ Two SPCN cables (FC 6001, FC 6006, FC 6007, FC 6008)
AIX operating system version	AIX 5L Version 5.1 or Version 5.2
Max. number of 7311-D20 I/O drawers attached to the p630	Two
Max. number of 7311-D20 I/O drawers attached to the p650	Eight

5.1.3 Configuration notes

The following are the configuration notes for the 731-D20 I/O drawer:

- ▶ Optional SCSI backplane (FC 6429), Ultra3 SCSI PCI Adapter (FC 6203) or Ultra3 SCSI RAID adapter (FC 2498), and cables (FC 4257) are required to support the disk drives installation.
- ▶ If the disk drives in the Model D20 are being used to support 2 LPARs, two SCSI adapters must be installed in the Model D20 (one to support each 6-pack of DASD).
- ▶ The optional SCSI backplane (FC 6429) consists of two 6-pack DASD bays.

5.1.4 6C4 I/O drawer attachment notes

The following are the I/O drawer attachment notes for the Model 6C4:

- ▶ The first Model D20 ordered requires two I/O cables and two SPCN cables to be on the order, with a second Model D20 requiring a third cable. These

cables can be ordered as part of the Model D20 order or as part of the Model 6C4 order.

- ▶ The 7028-6C4 must have FC 9575 or FC 6575 installed to attach the 7311-D20

Note: The configurator validates system connections when the remote I/O cables/SPCN cables are ordered on the Model D20. Anytime remote I/O cables/SPCN cables are ordered on the Model 6C4, system validation cannot be performed.

5.1.5 6M2 I/O drawer attachment notes

The following are the I/O drawer attachment notes for the p650:

- ▶ The p650 system allows up to eight I/O drawers to be attached to the system CEC drawer. I/O drawers should 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 6M2 drawer via the following cables:
 - RIO cables for data transfer
 - 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. p650 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.
- ▶ RIO ports are provided by remote I/O loop adapter cards. The first primary RIO adapter (FC 6411 or FC 6415) is required. An optional secondary RIO adapter (FC 6412 or FC 6416) provides two additional RIO ports. The first secondary RIO adapter must be installed prior to installing the second primary RIO adapter.
- ▶ Up to four RIO loops are available, each supporting up to four I/O drawers and up to a maximum of eight 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 in various lengths to attach I/O drawers within a single rack or across multiple racks, if desired.
- ▶ FC 3146 = 1.2 meters
- ▶ FC 3147 = 3.5 meters

- ▶ FC 3148 = 10 meters
- ▶ Power control for the I/O drawers is provided via one loop. The number of power control (SPCN) cables required is equal to one plus the number of I/O drawers attached to the system. A minimum of two power control cables (FC 6001 = 2 meters, FC 6006 = 3 meters, FC 6007 = 15 meters, or FC 6008 = 6 meters) are required for attachment of the first drawer. Each additional drawer requires one additional power control cable to complete the loop attachment.

Note: The configurator validates system connections when the RIO cables are ordered on the individual 7311-D10 or 7311-D20 drawers. If RIO cables are ordered on the 7038-6M2 server, system validation cannot be performed.

- ▶ The number of RIO or RIO-2 adapters must be less than or equal to the number of 2-way processor cards.
- ▶ Each RIO or RIO-2 adapter supports one RIO loop. Each RIO loop supports up to two 7311-D20 I/O drawers or up to four 7311-D10 I/O drawers.
- ▶ 7311-D10 and 7311-D20 I/O drawers are not supported in the same I/O loop.
- ▶ For systems with both 7311-D10 and 7311-D20 I/O drawers attached, the combined maximum number of I/O drawers is eight.
- ▶ To achieve the higher Remote I/O performance of RIO-2, the system must have a RIO-2 Enabled System Planar (FC 9581 or FC 6581). Also, to achieve RIO-2 performance, all adapters in a loop must be RIO-2. RIO and RIO-2 adapter combinations are supported in the same loop; however, the loop will operate at the standard RIO speed.

5.2 7311-D10 I/O Drawer

This section discusses the 7311-D10 I/O drawer.

Figure 5-6 on page 207 shows the 7311-D10 I/O drawer enclosure. Two D10 drawers can fit inside this enclosure.



Figure 5-6 7311-D10 I/O drawer enclosure

The IBM 7311 Model D10 I/O drawer is a new, rack-mountable expansion cabinet that can be attached to a pSeries 650 server. Each Model D10 drawer gives you six full length adapter slots. Up to eight Model D10 drawers are supported by the pSeries 650, with connections provided by Remote I/O adapters and cables.

The Model D10 requires 4U of vertical space in a 19-in. rack, such as the IBM 7014-T00 or 7014-T42. Two D10 drawers can fit side by side within the enclosure that provides rack mounting hardware.

The Model D10 offers a modular growth path for pSeries 650 systems with increasing I/O requirements. When a pSeries 650 is fully configured with eight attached Model D10 drawers, the combined system supports up to 55 PCI adapters.

5.2.1 Highlights

The Model D10 offers a modular I/O growth path for pSeries 650 systems. The characteristics are shown in Table 5-2.

Table 5-2 7311-D10 I/O drawer at a glance

Features	Detail
Max. number of 7311-D10 I/O drawers attached to the p650	Eight; provides up to 48 additional PCI slots for a p650
Rack-mount enclosure	4U rack-mount enclosure that holds one or two D10 drawers
PCI slots	Five hot-plug PCI-X slots and one hot-plug PCI slot supporting current and legacy 5V adapters

Features	Detail
Hardware minimum requirements	<ul style="list-style-type: none"> ▶ One Rack Mount I/O Drawer Enclosure, Two Position (FC 7311) ▶ Two AC Power Supply, 250W (FC 6278) ▶ One Remote I/O Loop Adapter: <ul style="list-style-type: none"> – (FC 6414) Remote I/O Loop Adapter or – (FC 6431) RIO-2 Remote I/O Loop Adapter (FC 6414) ▶ Two Remote I/O Cables (FC 3146, FC 3147, or FC 3148) ▶ Two Power Control Cables (FC 6001, FC 6006, FC 6007, or FC 6008) ▶ One AC Power Specify (FC 9172) ▶ One Language Group Specify (FC 9XXX) ▶ One Power Cord Specify (FC 99XX)
AIX operating system version	AIX 5L Version 5.1 or Version 5.2

5.2.2 I/O drawer attachment notes

The following are the I/O drawer attachment notes for the p650:

- ▶ The pSeries 650 system allows up to eight I/O drawers to be attached. It is recommended the I/O drawers be located in the same rack as the server drawer for service purposes; however, they can be mounted in separate racks, if desired.
- ▶ I/O drawers are connected to the 6M2 drawer via the following cables:
 - RIO cables for data transfer
 - 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. p650 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.
- ▶ Up to four RIO loops are available, each supporting up to four I/O drawers and up to a maximum of eight 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 in various lengths (FC 3146 = 1.2 meters, FC 3147 = 3.5 meters, and FC 3148 = 10 meters) to attach I/O drawers within a single rack or across multiple racks, if desired.

- ▶ Power control for the I/O drawers is provided from the p650 using one loop. The number of Power Control (SPCN) cables required is equal to one plus the number of I/O drawers attached to the system. A minimum of two Power Control cables (FC 6001 = 2 meters, FC 6006 = 3 meters, FC 6007 = 15 meters, or FC 6008 = 6 meters) is required for attachment of the first D10 drawer to the server. Each additional drawer requires one additional power control cable to complete the loop attachment.

5.2.3 Features

This section lists the internal features that can be added to the 7311-D20 and D10 I/O drawers. The status of a feature is indicative of these qualifications:

- X** Available. Indicates features that are available and orderable on the specified model.
- R** Must be removed. Indicates this feature is not supported and must be removed during a model conversion.
- W** Withdrawn. Indicates a feature that is no longer available to order.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, mice, language specify codes, and power cords, are not included.

The features of the 7311 D20 and D10 I/O drawers are listed in Table 5-3.

Table 5-3 Features of the 7311-D20 and D10 IO drawers

Feature	Description	7311 D10	7311 D20
Host attachment adapters			
2751	S/390 ESCON Channel PCI Adapter	X	
8398	SP Switch2 PCI-X Attachment Adapter		X
Asynchronous adapters			
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI bus	X	X
2944	128-Port Asynchronous Controller, PCI bus	X	X

Feature	Description	7311 D10	7311 D20
Asynchronous cables			
2936	Asynchronous Cable EIA-232/V.24	X	X
2934	Asynchronous Terminal/Printer Cable EIA-232	X	X
8131	128-Port Asynchronous Controller Cable, 4.5 Meter	X	X
8132	128-Port Asynchronous Controller Cable, 23cm (9-in.)	X	X
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232	X	X
8133	RJ-45 to DB-25 Converter Cable	X	X
8136	Rack Mountable Remote Asynchronous Node 16-Port EIA-232	X	X
SCSI adapters			
6204	PCI Universal Differential Ultra SCSI Adapter	X	X
6203	PCI Dual Channel Ultra3 SCSI Adapter	X	X
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter	X	X
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter		X
5711	PCI-X Dual Channel Ultra320 SCSI RAID Adapter (in cassette)	X	
5712	PCI-X Dual Channel Ultra320 SCSI Adapter		X
5710	PCI-X Dual Channel Ultra320 SCSI Adapter (in cassette)	X	
Serial adapters			
6231	128 MB DRAM Option Card	X	X
6230	Advanced SerialRAID Plus Adapter	X	X
6235	32 MB Fast-Write Cache Option Card	X	X
Fibre Channel adapters			
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus	W	W
6239	2 Gigabit Fibre Channel PCI-X Adapter	X	X
Graphics accelerators			
2848	POWER GXT135P Graphics Accelerator	X	X
2849	POWER GXT135P Graphics Accelerator with Digital Support	X	X

Feature	Description	7311 D10	7311 D20
Graphics cable			
4242	6 Foot Extender Cable for Displays (15 pin D-shell to 15 pin D-shell)	X	X
LAN adapters			
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	X	X
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	X	X
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	X	X
4959	IBM Token-Ring PCI Adapter	X	X
2946	Turboways 622 Mbps PCI MMF ATM Adapter	X	X
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	X	X
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	X	X
4962	10/100 Mbps Ethernet PCI Adapter II	X	X
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	X	X
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	X	X
WAN adapters			
2962	2-Port Multiprotocol PCI Adapter	X	X
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	X	X
6310	IBM ARTIC960RxD Quad Digital Trunk PCI Adapter		X
6312	Quad Digital Trunk Telephony PCI Adapter		X
6313	Quad Digital Trunk Telephony PCI Blind Swap Adapter	X	
WAN cables			
2861	ARTIC960Hx 4-Port EIA-232 Cable	X	X
2864	ARTIC960Hx 4-Port V.35 (DTE) Cable	X	X
2863	ARTIC960Hx 4-Port X.21 Cable	X	X
2710	ARTIC960Hx/RxD 4-Port E1 RJ-45 Cable	X	X
2709	ARTIC960Hx/RxD 4-Port T1 RJ-45 Cable	X	X
2873	IBM ARTIC960RxD Quad DTA, E1, 120 Ohm Balanced, 3M 4-Port Cable		X

Feature	Description	7311 D10	7311 D20
2874	IBM ARTIC960RxD Quad DTA, E1, 120 Ohm Balanced, 7.5M Extension Cable		X
2871	IBM ARTIC960RxD Quad DTA, T1, 100 Ohm, 3M 4-Port Cable		X
2872	IBM ARTIC960RxD Quad DTA, T1, 100 Ohm, 15M Extension Cable		X
2951	Cable, V.24/EIA-232	X	X
2952	Cable, V.35	X	X
2953	Cable, V.36/EIA-499	X	X
2954	Cable, X.21	X	X
Adapter mounting hardware			
4599	PCI Blind Swap Cassette Kit, Single Wide Adapters, Universal	X	
Cryptographic			
4963	PCI Cryptographic Coprocessor (FIPS-4)	X	X
4960	IBM e-business Cryptographic Accelerator	X	X
Miscellaneous internal system cards			
6413	RIO Ports to I/O Planar Riser Card		X
6417	RIO-2 Ports to I/O Planar Riser Card		X
6414	Remote I/O Loop Adapter	X	
6431	RIO-2 Remote I/O Loop Adapter	X	
SCSI internal cards			
4257	PCI Adapter to DASD Backplane Cable		X
SCSI external			
2424	0.6M 16-bit SCSI-2 System to System Cable	X	X
2425	2.5M 16-bit SCSI-2 System to System Cable	X	X
2118	Converter Cable, VHDCI to P, Mini-68 pin to 68 pin, 0.3M	X	X
2114	PCI SCSI Adapter 16-Bit Differential External Y Cable	X	X
Fibre Channel			
2456	LC-SC 50 Micron Fiber Converter Cable	X	X

Feature	Description	7311 D10	7311 D20
2459	LC-SC 62.5 Micron Fiber Converter Cable	X	X
Miscellaneous			
2877	IBM ARTIC960RxD Quad DTA, H.100, 4-Drop Cable	X	X
2733	IBM Long-wave Serial HIPPI PCI Adapter for RS/6000		W
2732	IBM Short-wave Serial HIPPI PCI Adapter for RS/6000	X	X
2737	Keyboard/Mouse Attachment Card - PCI	X	X
4353	H-100 Bus 8-position Cable	X	X
4692	Rack Status Beacon Cable, Junction Box Daisy Chain	X	X
4691	Rack Status Beacon Cable, Junction Box To Drawer Or Status Beacon	X	X
3124	Serial to Serial Port Cable for Drawer/Drawer	X	X
3125	Serial to Serial Port Cable for Rack/Rack	X	X
9004	Southern Hemisphere Designator for Monitors	X	X
Rack related			
3254	RIO Cable Support Brackets		X
6001	Power Control Cable (SPCN) - 2 meter	X	X
6006	Power Control Cable (SPCN) - 3 meter	X	X
6008	Power Control Cable (SPCN) - 6 meter	X	X
6007	Power Control Cable (SPCN) - 15 meter	X	X
3146	RIO-2 (Remote I/O-2) Cable, 1.2M	X	
3147	RIO-2 (Remote I/O-2) Cable, 3.5M	X	X
3148	RIO-2 (Remote I/O-2) Cable, 10M	X	X
SCSI disks			
3272	18.2 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly		W
3273	36.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly		X
3277	36.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly		X
3274	73.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly		X

Feature	Description	7311 D10	7311 D20
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly		X
3275	146.8 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly		X
Monitors			
3636	L200p Flat Panel Monitor	X	X
3635	T210 Flat-Panel Monitor	W	W
3628	IBM P260/P275 Color Monitor, Business Black, and Cable	X	X
3637	IBM T541H 15" TFT Color Monitor, Business Black, Captured Cable	X	X
Disk bays			
6429	Ultra320 SCSI Backplane for Hot-swap Disks		X
Routing indicators			
0986	CCS Customer Service Specify (US)	X	X
5001	Customer Service Specify	X	X
9006	Linux Ready Express Configuration Indicator Feature	X	
7311	Rack Mount I/O Drawer Enclosure, Two Position	X	
4693	Rack Status Beacon Junction Box	X	X
Rack ID			
4651	Rack Indicator, Rack #1	X	X
4652	Rack Indicator, Rack #2	X	X
4653	Rack Indicator, Rack #3	X	X
4654	Rack Indicator, Rack #4	X	X
4655	Rack Indicator, Rack #5	X	X
4656	Rack Indicator, Rack #6	X	X
4657	Rack Indicator, Rack #7	X	X
4658	Rack Indicator, Rack #8	X	X
4659	Rack Indicator, Rack #9	X	X



High-end servers: IBM @server pSeries 670 and 690

This chapter describes the following IBM @server pSeries high-end servers:

- ▶ 7040 pSeries 670
- ▶ 7040 pSeries 690
- ▶ Features of the 7040 pSeries 670/690

For additional information, refer to the redbook *IBM pSeries 670 and pSeries 690 System Handbook*, SG24-7040.

6.1 Positioning pSeries 670 and 690

The pSeries 690 is one of the most powerful UNIX servers ever built, with the best performance per processor in many application areas. The same holds for the pSeries 670 in the midrange. From copper-based, Silicon-on-Insulator (SOI) POWER4 microprocessor technology and multichip packaging expertise to the industry's most advanced self-management capabilities, the pSeries 670 and pSeries 690 are designed to provide the ultimate in performance, configuration flexibility, and availability.

The pSeries 670 and pSeries 690 are positioned to fulfill the applications' need for performance and scalability, and the reliability required in mission-critical environments. It is an ideal server on which to run corporate applications, such as Online Transaction Processing (OLTP), Enterprise Resource Planning (ERP), Business Intelligence (BI), and high-performance e-business infrastructures. It also provides attractive facilities for server consolidation with its flexible partitioning mechanism, as well as the dynamic logical partitioning function, possibilities, and advanced management functions.

Performance wise, the pSeries 670 and pSeries 690 excel in commercial processing, and the pSeries 690 set many new records in High-Performance Computing (HPC) applications. The balanced architecture, together with high-performance processors, offer unparalleled price/performance, helping to reduce costs not only on the server itself, but also on power, cooling, and software licenses.

In the case of very resource-consuming applications, or for consolidation of a very high number of servers, it is possible to group several pSeries 670 and pSeries 690 into clusters, (optionally with other pSeries, RS/6000, and SP servers), known as IBM @server Cluster 1600. Additional clustering information can be found at:

<http://www.ibm.com/servers/eservers/clusters/hardware/1600.html>

For more information on the AIX 5L operating system and related IBM products, the following link may be of interest:

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

For application availability on the AIX 5L operating system, the following link can be used for alphabetical listing and advanced search options, both for IBM software products and third-party software products:

<http://www.ibm.com/servers/aix/products/>

6.2 7040 pSeries technical overview

This section provides a description of the system enhancements and shows a picture of the modular design (Figure 6-1).




Figure 6-1 7040 modular design

The I/O drawers can contain up to 16 hot-swap capable disks structured into four 4-packs of 18.2 GB, 36.4 GB, 73.4 GB, or 146.8 GB disks. This delivers up to 2.34 TB maximum per I/O Drawer.

The drawer contains two planars with 10 slots per planar, 20 slots/drawer. The slots are hot-pluggable to permit PCI or PCI-X adapters to be added or replaced without extending the I/O drawer, all while the system remains available to the customer.

Complementing the power of the pSeries 670 or 690 is extraordinary self-managing capabilities. The system can detect faulty memory, processors, cache and PCI buses, and dynamically take them offline. It even makes the service call, all without requiring administrative action or interrupting operations. Growth is easy with hot-swappable disk drive and hot-plug PCI and PCI-X adapters. For the ultimate in availability, the systems can be clustered together with High Availability Cluster Multiprocessing software, the leader in UNIX availability solutions.

The pSeries 670 has some differences from the pSeries 690. Basically, it has the same technical advantages and Reliability, Availability, and Serviceability (RAS) features, but it is a single rack server and has several configuration limitations compared to the pSeries 690. Table 6-1 provides a quick overview of the differences between these two models.

Table 6-1 Differences between pSeries 670 and pSeries 690

Component	pSeries 690	pSeries 670
Number of CPUs/MCMs	8-32/1-4	4-16/1-2
4-way 1.1 GHz POWER4	No	Yes
8-way 1.1 GHz POWER4	Yes	Yes
8-way 1.3 GHz POWER4 Turbo	Yes	No
4-way 1.3 GHz POWER4 HPC	Yes	No
4-way 1.5 GHz POWER4+	No	Yes
8-way 1.5 GHz POWER4+	Yes	Yes
8-way 1.7 GHz POWER4+	Yes	No
Amount of memory (GB)	8-512	4-256
Number of memory cards	1-8	1-4
Number of I/O drawers	1-8	1/2 ^a - 3
Number of PCI slots	20-160	10-60
Internal battery feature	Yes	Yes
Number of racks	1-2	1
a. As an initial order, just one half I/O drawer can be used with 10 PCI slots.		

The MCMs used in the pSeries 670 are either 4- or 8-way modules. Since it is only possible to install two MCMs in the pSeries 670, only the inward-facing memory card slots can be used.

Note: It is possible to upgrade a pSeries 670 into a pSeries 690 by model conversion.

6.2.1 Central Electronics Complex (CEC 7040-671/681)

The Central Electronics Complex is a 17 EIA height drawer housing the processors and memory of the pSeries 670 and pSeries 690. The CEC contains the following components:

- ▶ The CEC backplane, where the components are mounted
- ▶ The multichip modules (MCMs), which contain the POWER4 processors
- ▶ Memory cards
- ▶ L3 cache modules
- ▶ I/O books, which provide the Remote I/O (RIO) ports for the I/O drawers, and the service processor function
- ▶ Fans and blowers for CEC cooling

Major design efforts have contributed to the development of the pSeries 670 and pSeries 690 to analyze single points of failure within the CEC to either eliminate them or to provide hardening capabilities to significantly reduce their probability of failure.

The front view of CEC is shown in Figure 6-2 on page 220. There are eight memory card slots available for pSeries 690 and four for pSeries 670. For detailed information about memory, see 6.2.2, “Memory subsystem” on page 228.




Figure 6-2 CEC front view

The rear view of CEC is shown in Figure 6-3 on page 221. The following slots are provided:

- ▶ Slots for distributed converter assembly (DCA) books and two capacitor books
 These slots are populated by up to six DCA books and up to two capacitor books. They supply electricity power to the CEC backplane and convert voltage.
- ▶ GX bus slots 0-3
 The GX bus slot 0 is used to insert the primary I/O book. The GX bus slots 1, 2, and 3 are used for the optional secondary I/O books. In Figure 6-3 on page 221, the slot 2 is populated with an IO book, while slots 1 and 3 are not. The pSeries 670 is restricted to at most one secondary I/O book, while the pSeries 690 can have from zero to three secondary books.




Figure 6-3 CEC rear view

In the center of the CEC, a CEC backplane is positioned vertically. As shown in Figure 6-4 on page 222, it provides mount spaces for up to four multichip modules (MCMs), sixteen level 3 (L3) cache modules, the eight memory cards, and the four I/O books. In the center position of the backplane is the clock card. It distributes sixteen paired clock signals to the four MCMs.




Figure 6-4 CEC backplane orthogonal view

POWER4 processor and MCM packaging

Several years ago, scientists working for microprocessor development set out to design a new microprocessor that would leverage IBM strengths across many different disciplines to deliver a server that would redefine what was meant by the term *server*. POWER4 is the result. It was developed by over 300 engineers in several IBM development laboratories. A new generation of this chip has been available since May 2003, which is called POWER4+. It has the same architecture as the initial POWER4 chip, but is able to sustain a higher clock frequency.

The pSeries 670 and pSeries 690 are based on these high-performance POWER4 chips. Up to 32 processors can be configured, and they are mounted together in multichip modules. This packaging methodology was carefully designed to provide increased levels of reliability by eliminating the tiered packaging levels of separate processor modules mounted on processor cards; these were mounted on a backplane. A pSeries 690 can have up to four MCMs installed.

The pSeries 670 is using the same high-performance POWER4 chips, but it supports up to 16 processors, which is a total of two MCMs installed.

Several versions of the POWER4 chips are available on pSeries 670 and pSeries 690 servers, with either 4 or 8 CPUs, executing at different speeds: 1.1GHz, 1.3 GHz, 1.5 GHz, or 1.7 GHz. All processors within pSeries 670 or

pSeries 690 server, must be identical and operate at the same speed. Table 6-2 describes supported combinations of processors.

Table 6-2 Supported combinations of processors

Architecture	Clock frequency	Number of processors per MCM	pSeries 670	pSeries 690	Designation
POWER4	1.1 GHz	4	4-way	N/A	Standard
		8	8-way or 16-way	8-way, 16-way, 24-way, or 32-way	Standard
	1.3 GHz	4	N/A	8-way or 16-way	HPC
		8	N/A	8-way, 16-way, 24-way, or 32-way	Turbo
POWER4+	1.5 GHz	4	4-way	N/A	
		8	8-way or 16-way	8-way, 16-way, 24-way, or 32-way	
	1.7 GHz	8	N/A	8-way, 16-way, 24-way, or 32-way	Turbo

The 1.1 GHz and 1.3 GHz 4-way MCMs are slightly different from the others. They use a POWER4 chip with a single core. The 1.3 GHz 4-way MCM (called the HPC feature) is specifically described in “Single core POWER4 processor feature” on page 227. It is a special option for technical applications that demand high memory bandwidth.

The chip has two processor cores and a level 2 (L2) cache, all in the same silicon substrate. The L2 cache is shared between the two cores through a crossbar switch known as the core interface unit (CIU). All these components are part of a single POWER4 chip, along with communication buses and a level 3 (L3) cache directory.

The following sections give a brief explanation about the POWER4 and MCM technologies.

For further detailed information about POWER4 and the memory subsystem in the pSeries 690, refer to the following whitepapers and redbook:

- ▶ *POWER4 System Microarchitecture*, found at:
<http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/power4.html>
- ▶ *IBM @server pSeries 690 Configuring for Performance*, found at:
http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_config.html
- ▶ *POWER4 Processor Introduction and Tuning Guide*, SG24-7041

POWER4 core

The internal micro-architecture of the POWER4 core is a speculative superscalar Out-of-Order execution design. Up to eight instructions can be issued each cycle, with a sustained completion rate of five instructions. Register rename pools and other Out-of-Order resources coupled with the pipeline structure allow the design to have over 200 instructions in flight at any given time.

In order to exploit instruction level parallelism, there are eight execution units, each capable of being issued an instruction each cycle. Two identical floating-point execution units, each capable of starting a fused multiply and add each cycle are provided. This combination can achieve four floating-point operations (FLOPs) per cycle per core. In order to feed the dual floating-point units, two load/store units, each capable of performing address generation arithmetic, are provided.

The microprocessor cores inside a POWER4 chip are a new design, although they support the same PowerPC instruction set architecture as prior pSeries processors. Each core has several execution units, along with level 1 (L1) instruction and data caches, and run at clock speeds of up to 1.7 GHz. Each L1 instruction cache is 64 KB and each L1 data cache is 32 KB. All data stored in the L1 data cache is available in the L2 cache, guaranteeing no data loss.

The initial POWER4 manufacturing process, known as *CMOS-8S3SOI*, implements seven-layer copper metallization, and 0.18 μm SOI CMOS technology. The POWER4+ manufacturing process is called *CMOS-9SSOI*, with a 0.13 μm SOI CMOS technology. Along with performance improvements, these technologies deliver high-reliability components that IBM considers to be fundamental for high-end, continuous operation servers. All the buses scale with processor speed.

L2 cache subsystem

The L2 cache on the POWER4 chip is composed of three separate cache controllers, connected to the two processor cores through a core interface unit. In POWER4 and POWER4+ chips, each L2 cache contains, respectively, 480 KB

and 512 KB, for a total of 1.44 MB per POWER4 chip, and 1.5 MB per POWER4+ chip.

Each L2 cache controller can operate concurrently and feed 32 bytes of data per cycle. The CIU connects each of the three L2 controllers to either the L1 data cache or the L1 instruction cache in either of the two processors.

Additionally, the CIU accepts stores from the processors across 8-byte wide buses and sequences them to the L2 controllers. Each processor has a non-cacheable (NC) unit associated with it, responsible for handling instruction serializing functions and performing any non-cacheable operations in the memory hierarchy. In a logical view, the NC units are part of the L2 cache.

The L2 cache on the POWER4 chip is dedicated to delivering data to the two cores as fast as they can process it, maximizing the overall performance. It has multiple ports and is capable of multiple concurrent operations. The L2 cache can provide data at the following peak rates:

- ▶ 1.1 GHz Standard system: 105.6 GBps to the two cores
- ▶ 1.3 GHz Turbo feature: 124.8 GBps shared to the two cores
- ▶ 1.3 GHz HPC feature: 83.2 GBps to the single core
- ▶ 1.5 GHz POWER4+ system: 144 GBps shared to the two cores
- ▶ 1.7 GHz POWER4+Turbo feature: 163.2 GBps shared to the two cores

L3 cache controller and directory

Each POWER4 processor chip, not each core, controls one L3 cache module. The L3 module itself is located outside the chip, but the L3 directory and L3 controller are located on the POWER4 chip. A separate functional unit, referred to as the Fabric Controller, is responsible for controlling data flow between the L2 and L3 controller for the chip and for POWER4 communication.

Each POWER4 chip has its own interface to the off chip L3 across two 16-byte wide buses, operating at one third of the processor frequency. To communicate with I/O devices, two 4-byte wide GX buses, operating at one third processor frequency, are used. Finally, each chip has its own Joint Test Action Group (JTAG) interface to the system service processor.

Also included on the chip are functions called pervasive functions. These include trace and debug facilities used for First Failure Data Capture (FFDC), Built-in Self Test (BIST) facilities, Performance Monitoring Unit, an interface to the service processor used to control the overall system, Power On Reset (POR) Sequencing logic, and error detection and logging circuitry.

Multichip modules

The POWER4 chips are packaged on a single module called multichip modules. Each MCM houses four chips (eight CPU cores) that are connected through chip-to-chip ports.

The chips are mounted on the MCM so that they are all rotated 90 degrees from one another, as shown in Figure 6-5. This arrangement minimizes the interconnect distances, which improves the speed of the inter-chip communication. There are separate communication buses between processors in the same MCM and processors in a different MCM (see Figure 6-6 on page 227).

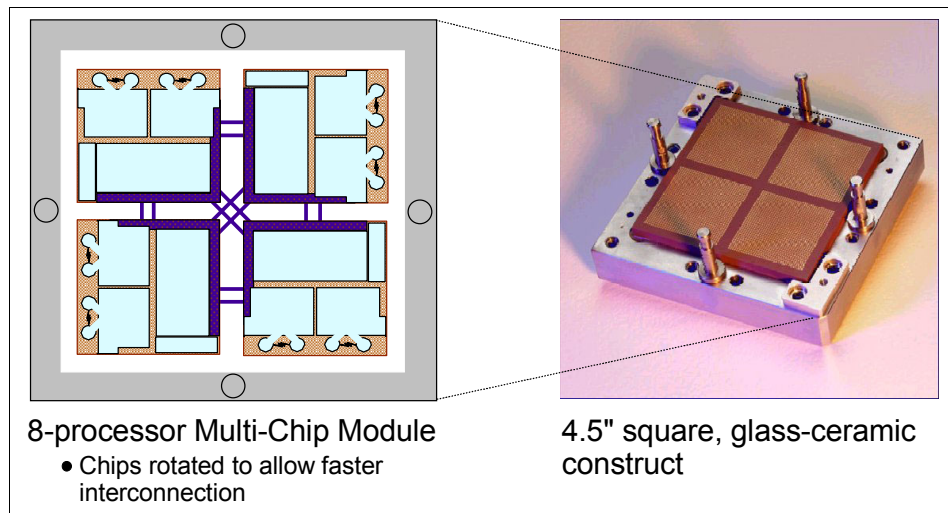


Figure 6-5 POWER4 multichip module

An internal representation of the MCM is shown in Figure 6-6 on page 227, with four interconnected POWER4 chips. Each installed MCM comes with 128 MB of L3 cache. This provides 32 MB of L3 cache per POWER4 chip. The system bus (L3 cache, GX Bus, and memory nest) operates at a 3:1 ratio with the processor frequency. Therefore the L3 cache to MCM connections operate at:

- ▶ 375 MHz for 1.1 GHz processors
- ▶ 433 MHz for 1.3 GHz processors
- ▶ 500 MHz for 1.5 GHz processors
- ▶ 567 MHz for 1.7 GHz processors

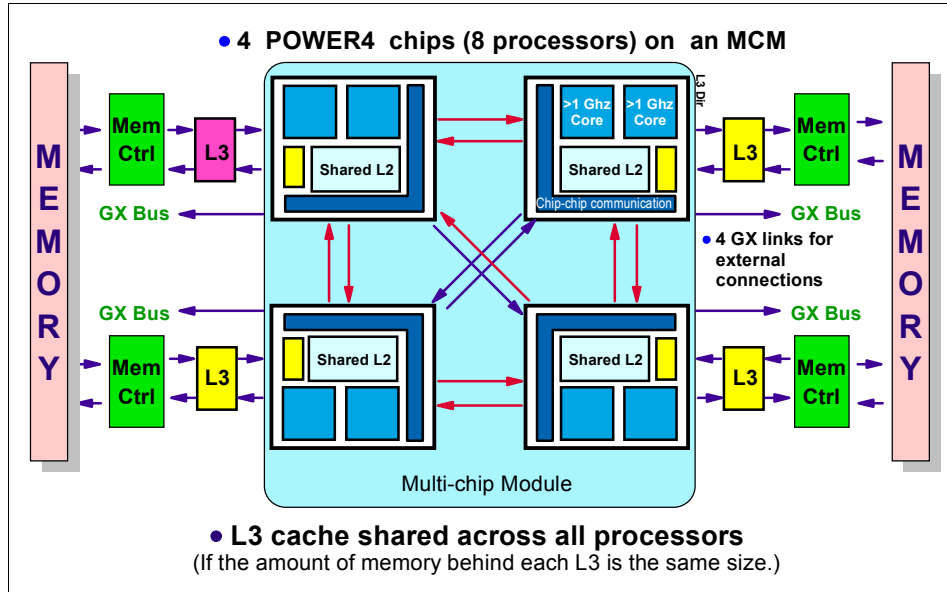


Figure 6-6 Multichip module with L2, L3, and memory

The MCM is a proven technology that IBM has been using for many years in mainframe systems (now IBM @server zSeries). It offers several benefits in mechanical design, manufacturing, and component reliability. IBM has also used MCM technology in the RS/6000 servers in the past. The IBM RS/6000 Model 580 was based on the POWER2 processor that has all its processing units and chip-to-chip wiring packaged in an MCM.

Single core POWER4 processor feature

As stated before, some technical applications benefit from very large bandwidth between processors and memory. The POWER4 processor delivers an exceptional bandwidth to the cores inside. For those applications that require extremely high bandwidth, the high-performance computing (HPC) feature is an attractive alternative. Instead of 8-way MCMs, you have 4-way MCMs with the same amount of L2 and L3 caches and the same bus interconnection (see Figure 6-7 on page 228).

This configuration provides twice the amount of L2 and L3 cache per processor and additional memory bandwidth, when compared to the pSeries 690 configured with 8-way processor MCMs. This additional cache and memory bandwidth available for each processor in this configuration may provide significantly higher performance per processor for certain engineering and technical environment applications.

Note: The HPC feature is only available on the pSeries 690 with a POWER4 1.3 GHz processor. It is not offered with the POWER4+ 1.5 or 1.7 GHz processors.




Figure 6-7 pSeries 690 1.3 GHz Turbo and 1.3 GHz HPC feature

6.2.2 Memory subsystem

This section mentions information related to the memory subsystem.

Figure 6-8 on page 229 shows the backplane of pSeries 670 and 690.

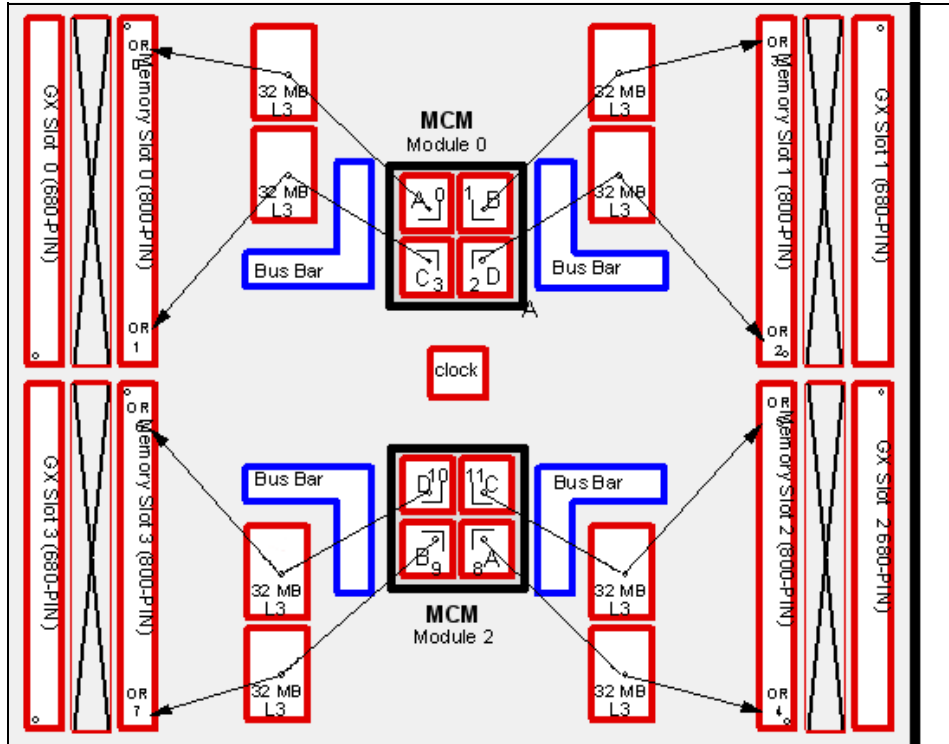


Figure 6-8 MCM, L3 cache, and memory slots relationship on backplane

The memory subsystem is physically composed of L3 cache modules and up to eight memory cards, both mounted on the CEC backplane. Since each L3 cache module is controlled by an L3 controller that physically resides in a POWER4 processor chip packaged in an MCM, there are tight relationships between MCMs, L3 cache modules, and memory cards.

MCM population order

The pSeries 690 can be configured with up to four MCMs mounted on the backplane, as shown in Figure 6-9 on page 230. The MCM 0 is always configured. The other optional MCMs are populated in this order: MCM 2, MCM 1, and MCM 3. In a two-MCM configuration (MCM 0 and MCM 2), because the spaces for MCM 1 and MCM 3 are not populated, you have to have a two-processor bus pass-through modules in this space to establish the link between two populated MCMs. In a three-MCM configuration (MCM 0, MCM 1, and MCM 2), since the space for MCM 3 is not populated, you have to have a processor bus pass through a module in this space.

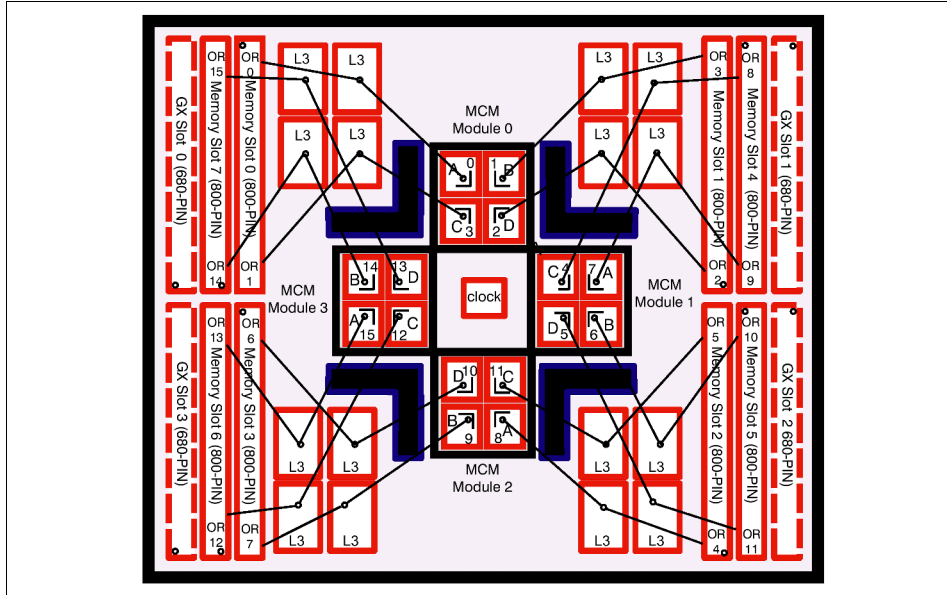


Figure 6-9 MCM, L3 cache, and memory slots relationship on backplane

I/O books

The I/O books plug into the GX slots in the CEC backplane and provide the Remote I/O (RIO) ports. The RIO ports are used to connect I/O drawers to the GX bus.

There are two types of I/O books available, called primary I/O book and secondary I/O book, as shown in Figure 6-10 on page 231. The primary I/O book is mandatory in all pSeries 670 and pSeries 690 and there must be exactly one per system. The secondary I/O books are optional. Their number depends on the requirements for external I/O drawer attachments. A pSeries 670 can contain at most one secondary I/O book, while a pSeries 690 can contain up to three secondary I/O books.

- ▶ The primary I/O book
 - Contains the following ports:
 - Two HMC connections ports: HMC1 and HMC2
 - Two native serial ports: S1 and S2

These two native serial ports are under the control of the service processor and are not available for functions, such as HACMP heartbeat cabling or UPS control, which require fully dedicated ports.

- An operator panel port

- A diskette drive port
- A bulk power controller (BPC) Y-cable connector
 - This port is used to connect between two BPCs on the bulk power assembly (BPA) and the service processor using a Y-cable.
- Contains the service processor.
- Contains non-volatile random access memory (NVRAM).
- Contains four RIO ports for the first I/O drawer (mandatory) and an optional second I/O drawer.
- This I/O book plugs into the GX slot 0 in the CEC backplane.
- ▶ Each secondary I/O book
 - Contains eight RIO ports.
 - Plugs into the GX slots of the CEC backplane in this order: GX slot 2, GX slot 3 and GX slot 1. In Figure 6-10, the RIO Port numbers correspond to the first secondary I/O book. For the second and third secondary books, the port number would be, respectively, 6, 7, 12, and 13, and 2, 3, 8, and 9.




Figure 6-10 Primary and secondary I/O books

There are two sets of I/O books, which differ by the technology used in the base card, a riser, and a daughter card installed in the book. The first set is referred to as RIO while the second is called RIO-2.

- ▶ The RIO primary and secondary books are:
 - Primary: FC 6404, Service Processor and RIO Loop Attachment, Two Loops
 - Secondary: FC 6410, RIO Loop Adapter, Four Loops
- ▶ The RIO-2 books are called:
 - Primary: FC 6418, Service Processor and RIO-2 Loop Attachment, Two Loops
 - Secondary: FC 6419, RIO-2 Loop Adapter, Four Loops

The RIO-2 books provide improvement in I/O performance and throughput required by the new 1.5 and 1.7 GHz processors. The RIO-2 primary book contains a larger NVRAM than the RIO primary book, and provide the increased support from 16 to 32 LPARs. There are restrictions in the use of the books:

- ▶ All I/O books must be of the same technology. You cannot mix FC 6404 and 6419 or FC 6418 and FC 6410 in the same system.
- ▶ The RIO books can only be used with the 1.1 and 1.3 GHz processors.
- ▶ The RIO-2 books can be used with the 1.1, 1.3, 1.5, and 1.7 GHz processors.
- ▶ When using RIO books, the maximum number of supported LPARs is 16.
- ▶ With RIO-2 books, it is possible to instantiate up to 32 LPARs, even with 1.1 GHz and 1.3 GHz systems.

6.2.3 I/O drawer - (7040-61D)

The 7040-61D I/O drawer provides PCI adapter slots and internal disk capabilities for use with pSeries servers such as the pSeries 670 (7040-671) or pSeries 690 (7040-681). The Model 61D is a 4U drawer that mounts in the 24-in. 7040-61R System Rack and connects to the system Central Electronics Complex using remote I/O cables.

Each Model 61D I/O drawer provides twenty blind swap PCI slots and sixteen hot-swap disk bays. The Model 61D utilizes redundant power converters and power cabling to ensure high reliability and availability. It utilizes 350V bulk power supplied from the 7040-61R system rack.

Note: The pSeries 670 supports up to three I/O drawers, while the pSeries 690 supports up to eight I/O drawers. The pSeries 670 also supports a half I/O drawer configuration (10 PCI slots).

Figure 6-11 shows the rear view of an I/O drawer, with the PCI slots and riser cards that connect to the RIO ports in the I/O books.




Figure 6-11 I/O drawer rear view

As for I/O books, the I/O drawers exist in two technologies: RIO and RIO-2. I/O drawers of both technologies offer the same number of PCI slots, the same number of disks, and are packaged in the same chassis. The difference is the type of I/O planar that is installed inside the drawer. Externally, the only visible difference between the two technologies is the shape of the cable connectors on the RIO Riser card (see Figure 6-12 on page 234).

- ▶ RIO connectors have a thumbscrew retention physical connector.
- ▶ RIO-2 connectors have a bayonet retention physical connector.




Figure 6-12 Difference between RIO and RIO-2 connectors

The I/O Drawer for pSeries 670 and pSeries 690 has its own product number: 7040-61D, which refers to the two technologies. For ordering purposes, the difference between them is the Feature Code of the configured I/O planar:

- ▶ RIO drawer uses the FC 6563, I/O Drawer PCI Planar, 10 slots, 2 Integrated Ultra3 SCSI Ports.
- ▶ RIO-2 drawer is configured with the FC 6571, I/O Drawer PCI-X Planar, 10 slots, 2 Integrated Ultra3 SCSI Ports.

A logical connection view of an RIO I/O drawer is shown in Figure 6-13.

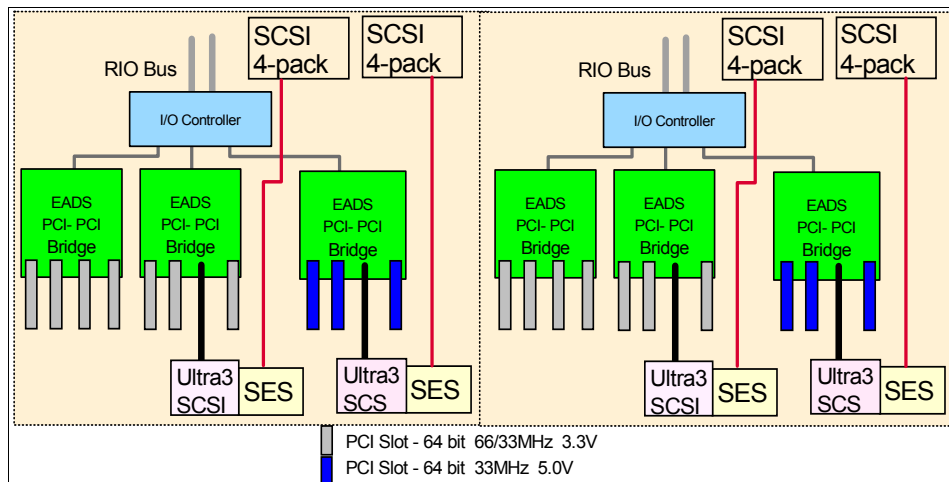


Figure 6-13 Logical view of an RIO drawer

For a RIO-2 I/O drawer, the topology of the logical connection view is identical, as shown in Figure 6-14 on page 235. The PCI bridge uses an EADS-X chip instead of an EADS chip. The PCI slots' differences are explained later in this section.




Figure 6-14 Logical view of an RIO-2 drawer

Each of the 4-packs supports up to four hot-swappable Ultra3 SCSI disk drives, which can be used for installation of the operating system or storing data. There are three different disk capacities: 18.2 GB, 36.4 GB, 73.4 GB, and 146.8 GB, and they have the following characteristics:

- ▶ Form factor: 3.5-in., 1-in. (25 mm) high
- ▶ SCSI interface: SCSI Ultra3 (fast 80) 16 bit
- ▶ Rotational speed: 10,000 RPM and 15,000 RPM

Optionally, external disk subsystems can be connected to serve as operating system disks. The pSeries 670 and pSeries 690 and their logical partitions can boot¹ from SCSI, SSA, and Fibre Channel disk subsystems.

The RIO riser cards are connected to the planar boards. The RIO ports of each riser card are connected through IO loops to RIO ports on I/O books in the CEC. The PCI slots have different characteristics in the RIO and RIO-2 drawers:

▶ RIO Drawer

On each planar board, the first seven PCI slots have a 3.3V PCI bus signaling and operating at 66 MHz or 33 MHz, depending on the adapter. The last three PCI slots have a 5V PCI bus signaling and operating at 33 MHz. All PCI slots are PCI 2.2 compliant and are Hot-Plug enabled, which allows most PCI adapters to be removed, added, or replaced without powering down the system. This function enhances system availability and serviceability.

¹ The boot capability from different disk technology depends on the operating system support. AIX supports these types of disk for a boot device.

▶ RIO-2 Drawer

On each planar board, the ten PCI-X slots have a 3.3V PCI bus signaling and operating at 33 MHz, 66 MHz, or 133 MHz, depending on the adapter. All PCI slots are PCI 2.2 compliant and are hot-plug enabled.

PCI adapters have different bandwidth requirements, and there are functional limitations on the number of adapters of a given type in an I/O drawer or a system. The limitations differ for the RIO and RIO-2 drawers.

The limitations are of several types:

- ▶ Maximum number of adapters of one type per IO planar
- ▶ Maximum number of adapters of one type per IO drawer
- ▶ Maximum number of adapters of one type per LPAR
- ▶ Maximum number of adapters of one type per pSeries 690 or per pSeries 670
- ▶ Maximum combination of specific adapters

The complete set of limitations are described in the *PCI Adapter Placement References*, SA38-0538. This book is regularly updated and should be considered as the reference for any questions related to PCI limitations.

6.2.4 Media drawer

The media drawer provided in the base configuration contains an operator panel, a 1.44 MB diskette drive, and two sets of two removable media bays. The operator panel and the diskette drive are each connected to the service processor book in the CEC using point-to-point cables, and are powered from the CEC standby power. Each set of media devices is powered and logically driven by an I/O planar in the first I/O drawer, as shown in Figure 6-15 on page 237.

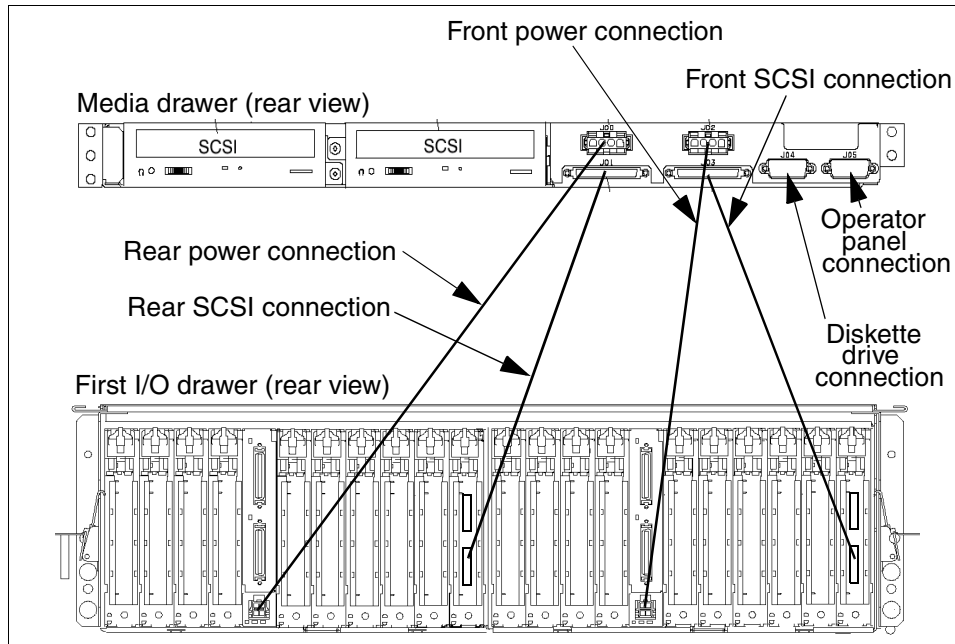


Figure 6-15 Media drawer power and SCSI connection

The media drawer is split into two separate sections. The front section houses the system operator panel, a diskette drive, and two media bays. The rear section provides space for two additional media devices. The front and rear sections of the media drawer are in two separate SCSI buses, and must be connected to separate SCSI PCI adapters on the I/O drawers. Therefore, these devices configured in the front and rear bays of the media drawer can either be in the same partition or in different partitions.

6.2.5 System Rack - (7040-61R)

The 7040-61R System Rack provides space for mounting system drawer components for pSeries servers such as the pSeries 670 (7040-671) or pSeries 690 Central Electronics Complex (7040-681) and the 7040-61D I/O drawers. The Model 61R is a 24-in. rack that provides 42U of rack space.

The 7040-61R incorporates a 350V DC bulk power subsystem to provide power for the components within the rack. The bulk power subsystem incorporates redundant bulk power assemblies mounted in the front and rear sections of the top 8U of the rack.

Power subsystem

The power subsystems on the pSeries 670 and pSeries 690 are identical and provide full redundancy. The power subsystem consists of redundant bulk power assemblies, bulk power regulators, power controllers, power distribution assemblies, DC power converters, and associated cabling.

The bulk power assembly (BPA) is the main power distribution unit for the pSeries 670 and pSeries 690. The redundant bulk power assembly converts AC input to DC and distributes power at 350V to each drawer where conversion is made to the required chip level. It is composed of several power converters and regulators. Two BPAs are mounted in front and rear positions and occupy the top of the rack as an eight EIA height drawer unit (see Figure 6-16). They are powered from separate power sources using separate line cords.

Each BPA consists of a stack of components (from bottom to top):

- ▶ Three Bulk Power Regulator (BPR)
- ▶ One Bulk Power Controller (BPC)
- ▶ Three Bulk Power Distributor (BPD). The first BPD must be installed on all systems. The second and third are optional, depending on the number of I/O drawers installed in the system.

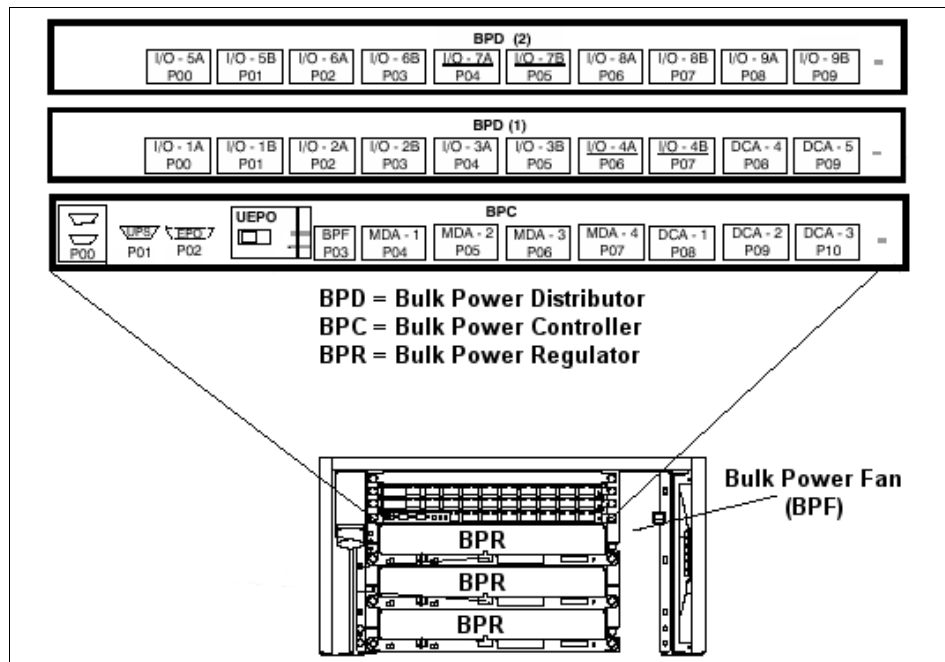


Figure 6-16 Power subsystem locations in BPA

The BPA distributes power to the I/O drawers and the DCA in the CEC using appropriate power cables, plugged into the BPC and BPD connectors (The number of the DCA in the CEC depends on the number of installed MCMs and I/O books).

Constant power monitoring hardware assists in the detection of failures in the power subsystem that require maintenance. The power supplies, power converters, and cables can be serviced without system disruption (see Figure 6-16 on page 238).

Internal battery feature

The internal battery feature (IBF) provides emergency power in case of a power outage. The pSeries 670 and pSeries 690 incorporate an Early Power Off Warning (EPOW) capability that assists in performing an orderly system shutdown in the event of a sudden power loss. The IBF protects against power line disturbances and provides sufficient power to enable the system to shut down gracefully in the event that the power sources fail.

Up to two IBF enclosures can be installed in the base rack configuration. One additional IBF enclosure can be installed in an expansion rack.

The hold-up time of the IBF will vary considerably with conditions present at the time of the power loss. Ambient temperature, age, and past use of the battery can affect the discharge time.

For more detailed information about the IBF hold-up times for the pSeries 670 and pSeries 690, see the 7040 sections in Chapter 1, “Physical Characteristics of Systems”, in *Site and Hardware Planning Information*, SA38-0508.

Attention: The document *Site and Hardware Planning Information*, SA38-0508, is updated very often. The version of this document that includes information related to the related to the pSeries 670 and pSeries 690 models announced in May 2003 is version 18. If you have an old version of this document, please download the latest version from the following IBM Web page:

http://www.ibm.com/servers/eserver/pseries/library/hardware_docs/index.html

Cooling

Several fans and blowers provide air cooling for the pSeries 670 and pSeries 690. The power supplies, the CEC, and each I/O drawer has a group of redundant cooling devices. They operate at normal speed when all the cooling devices are operational, and increase their speed if one device stops working in order to maintain the constant air flow and provide proper cooling.

The CEC is cooled by the following devices:

- ▶ Four high-pressure, high-flow blowers in the base of the CPU cage that pull air down the front of the cage through the memory cards and exhaust air up the rear of the cage through the power converters and RIO books.
- ▶ Two additional identical blowers in the front of the cage blow air directly down on the MCMs and L3 modules, which is then exhausted out the top rear of the cage.

Each fan or blower is referred to as an Air Moving Device, which consists of a Motor/Scroll Assembly (MSA) and a Motor Drive Assembly (MDA).

- ▶ The MSA consists only of a motor, wheel, and metal housing.
- ▶ The MDA has two 350V and communication inputs (one from each BPA) and a three-phase output designed to drive a brush-less DC motor using back EMF commutation control. The power and RAS infrastructure is used to control and activate the Air Moving Devices through the communication ports.

A micro controller within the MDA monitors and controls the motor drive and provides precise speed control capability. Each Air Moving Device can draw 15-300 watts of power, depending on speed setting and back pressure, with the maximum power for a set of two blowers being 500 watts. The Air Moving Device design is identical to that of the Air Moving Device used in the zSeries servers. All Air Moving Devices from the CEC can be removed and replaced without powering down the system.

Please note that each I/O drawer has the cooling devices that have the same functionality within the CEC, and also offers N+1 redundancy.

6.2.6 IBM Hardware Management Console for pSeries

The IBM Hardware Management Console for pSeries (HMC) provides a set of functions that are necessary to manage the pSeries 670/690 LPAR configurations.

The HMC also provides a service focal point for the systems it controls. It is connected to the support (service) processor of the system by a dedicated serial link. The HMC provides tools for problem determination and service support, such as call-home and error log notification through an analog phone line.

The HMC is a dedicated function device. It is utilized only for the control and service functions of the pSeries servers it serves. It is not available for use as a general purpose computing resource.

The HMC is packaged in either a desktop or 19-in. rack-mount version.

Figure 6-17 shows an example of the graphical user interface on the HMC.




Figure 6-17 Graphical user interface on the HMC

For further information about how to use and manage HMC, refer to the following publications:

- ▶ *IBM Hardware Management Console for pSeries Installation and Operations Guide*, SA38-0590
- ▶ *IBM Hardware Management Console for pSeries Maintenance Guide*, SA38-0603

Note: The HMC is a mandatory feature of the pSeries 670 and pSeries 690.

6.2.7 RAS features

Both the pSeries 670 and pSeries 690 bring new levels of availability and serviceability to the enterprise-level UNIX servers. Advanced functions presented on the previous models were enhanced, and new technologies have been added, to reduce faults and minimize their impacts. High-quality components and rigorous verification and testing during the manufacturing processes contribute to the reduction of failures in the system. The diagnostic

goal of these RAS features is to isolate the Field Replaceable Units (FRU) callouts to 95 percent to a single FRU; there will still be a chance that more than one FRU is failing, but the implemented features in the pSeries 670 and pSeries 690 will keep the customer impact as small as possible.

The pSeries 670 and pSeries 690 RAS design enhancements can be grouped into four main areas:

- ▶ Predictive functions: These are targeted to monitor the system for possible failures, and take proactive measures to avoid the failures.
- ▶ Redundancy in components: Duplicate components and data paths to prevent single points of failure.
- ▶ Fault recovery: Provide mechanisms to dynamically recover from failures, without system outage. This includes dynamic deallocation of components and hot-swap capability.
- ▶ Serviceability features: Enable the system to automatically call for support, and provide tools for rapid identification of problems.

6.3 7040-671 IBM @server pSeries 670

This section discusses the pSeries 670 Model 671. Figure 6-18 show the pSeries 670.



Figure 6-18 Front view pSeries 670

6.3.1 Highlights of pSeries 670

This sections explain the highlights of the pSeries 670 system.

- ▶ IBM SOI/copper chip technology.
- ▶ Mainframe-inspired LPAR and self-managing capabilities.
- ▶ Cost-efficient growth path to the future.
- ▶ DDR System memory, expandable from 4 GB to 256 GB.
- ▶ IBM @server Capacity Upgrade on Demand for Processors.
- ▶ IBM @server Capacity Upgrade on Demand for Memory.
- ▶ Up to three 7040-61D I/O drawers per server; each I/O drawer supports 20 blind-swap PCI or PCI-X bus slots and up to 16 Ultra3 SCSI disks that can be hot-plugged.
- ▶ A redundant system power that provides a modular design that allows adding power regulators and power converters as required based on the system configuration.
- ▶ Industry recognized AIX 5L operating system license is included with each system.
- ▶ Hardware split into as many as sixteen LPARs, each functioning as a computer within a computer with its instance of the operating system.
- ▶ Excellent reliability, availability, and serviceability (RAS) providing:
 - Dynamic processor deallocation
 - Error checking and correcting (ECC) system memory and cache
 - I/O link failure recovery
 - Disk drives and PCI cards that can be hot-swapped
 - Environmental sensing
 - Redundant power and cooling
 - Service processor for integrated system monitoring
 - Concurrent diagnostics
- ▶ 64-bit system scalability
 - 4-way, 8-way, or 16-way packaged on two Multi-Chip Modules (four or eight POWER4 processors per MCM)
 - 1.1 GHz POWER4 processors
 - 1.5 GHz POWER4+ processors
 - 256 MB of L3 ECC cache per MCM

6.3.2 7040 pSeries 670 minimum and standard features

This section describes the system minimum, standard features of a pSeries 670 (see Table 6-3).

Table 6-3 pSeries 670 minimum and standard feature

Specification	Detail
Processor	4-way SMP (one 4-way MCM); 1.5GHz POWER4+ or 1.1GHz POWER4
L3 cache	128 MB
Memory (minimum)	4 GB
Internal disk drives	Two 36.4 GB Ultra SCSI
Disk bays	16 hot-swappable using one 7040-61D I/O drawer
Media bays	Five (four available)
Expansion slots	20 PCI or PCI-X (64-bit) using one 7040-61D I/O drawer
PCI bus width	32- and 64-bit
I/O adapters	Two integrated Ultra3 SCSI controllers
Standard Ports	Two serial ports for connecting Hardware Management Console for pSeries
Operating systems	AIX 5L Versions 5.1/5.2 SuSE Linux Enterprise Server Version 8 (runs in LPAR only)
Software Requirements	AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package (APAR IY 28102), or later, or AIX 5L Version 5.2, or later.

6.3.3 7040 pSeries 670 system expansion

Table 6-4 shows the possible maximum processor, memory, and storage configurations.

Table 6-4 pSeries 670 system expansion

System expansion	
SMP configuration	8- or 16-way SMP (one or two 8-way MCMs); 1.5GHz POWER4+ or 1.1GHz POWER4
L3 cache	128 MB per MCM (256 MB maximum)

System expansion	
Memory	Up to 256 GB (ECC, Chipkill)
PCI expansion slots	Up to 60 adapters using two additional 7040-61D I/O drawers
Disk bay expansion	Up to 32 hot-swappable disk bays using two additional 7040-61D I/O drawers; up to 7.0 TB (36.4 GB, 73.4 GB, and 146.8 GB disk drives available)
Optional battery backup	Up to two
Attachment	SP System Attachment Adapter for use in a Cluster 1600 configuration

6.3.4 Configuration notes

In the following section, we describe some configuration notes about the pSeries 670 system.

pSeries 670 Central Electronics Complex (CEC)

This section provides information about the limitations of the pSeries 670 CEC.

- ▶ The p670 is powered by one or two multi-chip processor modules. Each Multi-Chip Module (MCM) contains either four or eight processors. The p670 is available in 4-, 8-, or 16-way configurations of 1.1 GHz POWER4 processors or 1.5 GHz POWER4+ processors (4 or 8 processors per MCM). The following configurations are supported:
 - One 4-way POWER4 or POWER4+ MCM with 128 MB L3 Cache
 - One 8-way POWER4 or POWER4+ MCM with 128 MB L3 Cache
 - Two 8-way POWER4 or POWER4+ MCMs with 128 MB L3 Cache
- ▶ All processors in a p670 must operate at the same speed.
- ▶ Each processor contains 32 KB of data cache and 64 KB of instruction cache.
- ▶ Each processor MCM contains 5.6 MB of Level 2 cache, which is shared between the processors. This provides 1.44 MB of L2 cache per POWER processor for four processor MCMs and 0.72 MB of L2 cache per processor for eight processor MCMs. POWER4+ processors have a 1.5 MB cache.
- ▶ Each installed processor MCM is supported with 128 MB of Level 3 cache.
- ▶ A programmable processor clock card (FC 5251) is required to provide system and processor clocking.
- ▶ p670 processor MCMs require one capacitor book (FC 6198) to operate.

- ▶ p670 servers with 1.5 GHz POWER4+ processors must be equipped with feature Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418).
- ▶ The p670 utilizes redundant power throughout its design. It implements redundant bulk power assemblies, bulk power regulators, power controllers, power distribution assemblies, DC power converters, and associated cabling.
- ▶ Power for the p670 CEC is supplied from DC bulk power assemblies in the 7040-61R rack. The bulk power is converted to the power levels required for the CEC using DC power converters. Two DC power converters (FC 6170) and attachment cables (FC 6181 and 6161) are required for the CEC and the first processor MCM, memory, and miscellaneous CEC components. The addition of a second MCM requires an additional DC power converter (FC 6189) and cable group (FC 6182).
- ▶ One Media Drawer (FC 8692) is required for each p670 server. The media drawer is a 1U rack drawer that is located directly below the CEC at the 17 U location of the 7040-61R rack. It incorporates the system operator panel, a diskette drive, and up to four media devices.
- ▶ The media drawer is split into two separate sections. The front section houses the system operator panel, a diskette drive, and two media bays. Utilization of the rear section is optional to provide space for two additional media devices if desired.
- ▶ The front and rear sections of the media drawer must be connected to separate ports on SCSI PCI adapters in the first 7040-61D I/O drawer, which is located in drawer position FC 4609. When connecting to a SCSI adapter with a 68 pin P-style connector, the I/O drawer to media drawer SCSI cable FC 2122 is used. When connecting to a SCSI adapter with a Mini-68 pin, VHDCI connector converter cable FC 2118 is used in conjunction with FC 2122. The SCSI adapters and cables for this attachment are ordered with the first 7040-61D I/O drawer.
- ▶ Power for each half of the media drawers is provided by the first I/O drawer using one or two I/O Drawer to Media Drawer Power Cables (FC 6179) ordered with the 7040-61D I/O drawer.
- ▶ One media device capable of reading CD-ROM media is required for each p670 system.
- ▶ The 4-mm tape drive (FC 6158) is allowed in the rear bays of the media drawer only if the operating environment is maintained at 24° C (75.2° F) or below and only on systems installed at 2,134 meters (7,000 ft.) altitude or below.
- ▶ The DVD-RAM drive (FC 2623), CD-ROM drive (FC 2624), and DVD-ROM drive (FC 2634) are limited to systems without Primary Integrated Battery

Backup (FC 6200) or Redundant Integrated Battery Backup (FC 6201) when installed in the Media Drawer (FC 8692) rear bay locations.

- ▶ The IBM 80/160 GB Internal Tape Drive with VXA Technology (FC 6120) is limited to systems without Primary Integrated Battery Backup (FC 6200) or Redundant Integrated Battery Backup (FC 6201) installed, Media Drawer (FC 8692) front bay locations only, and a maximum system ambient operating temperature of 28° C (82.4 F) at a maximum operating altitude of 2,134 m (7,000 ft.). Lower altitude have higher maximum ambient operating temperatures. Refer to the *7040 System Planning Guide* (found at http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm) for additional details.
- ▶ USB keyboards supporting various language groups can be ordered with the 7040-61D I/O Drawer for use with native-attached displays on the pSeries 670. A three-button USB mouse connects to the USB keyboard.

pSeries 670 memory

This section provides information about the limitations of the pSeries 670 memory.

- ▶ The pSeries 670 has four memory slots. These memory slots utilize inward facing memory cards.
- ▶ The number of active memory positions on a p670 system is related to the number of populated MCM locations. Two memory positions are activated for each MCM installed.
- ▶ The first MCM installed activates the first and second inward facing memory positions.
- ▶ The second MCM installed activates the third and fourth inward facing memory positions.
- ▶ Minimum system memory is 4 GB. Maximum system memory is 256 GB.
- ▶ Memory is available in 4 GB, 8 GB, 16 GB, 32 GB, and 64 GB increments.
- ▶ Memory is available in 500 MHz and 567 MHz speeds.
- ▶ 500 MHz and 567 MHz memory can be simultaneously installed in a 7040-671 server. In this configuration, all memory will perform at 500 MHz.
- ▶ 500 MHz memory allows maximum memory subsystem performance for 7040-671 servers with POWER4 or POWER4+ processor MCMs operating at speeds up to 1.5 GHz.
- ▶ IBM recommends systems configured with 4 GB of memory or greater to have access to a 4-mm or 8-mm tape drive for submission of system dump information, if required. This function may be accomplished using locally attached or network attached devices as appropriate.

- ▶ The following memory configuration considerations are recommended to ensure the best performance is obtained for the pSeries 670 system:
 - Memory should be installed in identical pairs.
 - All activated memory locations should be populated with memory cards.
 - Memory sizes should be balanced as closely as possible across all populated MCM locations.

Table 6-5 shows the supported memory card configurations on pSeries 670.

Table 6-5 Supported memory configurations for pSeries 670

MCMs installed	Memory size in GB	Memory slots 0 & 1	Memory slots 2 & 3
1	4	4 + 0	N/A
	8	4 + 4	
	8 ^a	8 + 0	
	16	8 + 8	
	32	16 + 16	
	64	32 + 32	
	128	64+64	
2	8	4 + 4	0 + 0
	16	4 + 4	4 + 4
	16 ^a	8 + 8	0 + 0
	24	8 + 8	4 + 4
	32	8 + 8	8 + 8
	48	16 + 16	8 + 8
	64	16 + 16	16 + 16
	96	32 + 32	16 + 16
	128	32 + 32	32 + 32
	192	64+64	32 + 32
	256	64+64	64+64

a. This is a valid configuration, but not a recommended configuration.

IBM Hardware Management Console for pSeries (HMC)

This section provides information about the limitations of the pSeries 670 HMC.

- ▶ The following Hardware Management Consoles can be used to control a p690 server:
 - HMC FC 7315 and FC 7316 purchased with a 7040-671 or 7040-681 server (withdrawn from marketing)
 - IBM 7315-C01 Hardware Management Console (withdrawn from marketing)
 - IBM 7315-C02 Hardware Management Console
 - IBM 7315-CR2 rack-mounted Hardware Management Console
- ▶ Each HMC must be configured with a graphics display and a keyboard. The 7315-CR2 rack-mounted HMC can be connected to the IBM 7316-TF2 rack-mounted display and associated VGA switch. Several HMCs can, therefore, share the same display, mouse, and keyboard. Various native language keyboards are available for the HMC. However, the HMC displays menus and control information in English regardless of the keyboard utilized.
- ▶ The number of servers an HMC can control is dependent on the particular HMC and the server environment.
- ▶ HMC FC 7315 (ordered with a 7040-681 server) can control up to four pSeries servers in a non-clustered environment.
- ▶ If redundant HMC function is desired, the pSeries 670 can be attached to two separate hardware management consoles.
- ▶ An Ethernet connection between the HMC and each active partition on the pSeries server is recommended for all installations and required for dynamic LPAR configurations. The HMC provides two integrated serial ports. One serial port is required for each pSeries server controlled by the HMC. An additional serial port is required for modem attachment if the Service Agent call-home function is implemented. When more than two serial ports are required, 8-port (FC 2943) and 128-port (FC 2944) asynchronous adapters should be utilized.

Note: 128-port asynchronous adapters are not supported on the FC 7316 HMC.

- ▶ Up to two 8-port or 128-port asynchronous adapters are allowed per HMC.
- ▶ 16 pSeries servers and 64 LPARS in a mixed server environment. A mixed server environment can contain a combined maximum of eight p670 and/or p690 servers.

- ▶ The 7315-C02 or 7315-C01 or FC 7316 HMC (ordered with a 7040-671 or 681 server) can control the following numbers of servers in a non-clustered environment:
 - 12 pSeries 670 and/or p690 servers with up to 64 LPARS or
 - 16 pSeries 655 servers with up to 32 LPARS or
 - 16 pSeries 630 or pSeries 650 servers with up to 64 LPARS or
 - 16 pSeries servers and 64 LPARS in a mixed server environment. A mixed server environment can contain a combined maximum of eight p670 and/or p690 servers.

Logical partitioning (LPAR)

This section provides information about the limitations of logical partitioning of the pSeries 670.

- ▶ The pSeries 670 can be divided into as many as 16 logical partitions. System resources can be dedicated to each LPAR.
- ▶ p670 servers configured with Support Processor with Remote I/O Loop Attachment (FC 6404) can support up to 16 LPARs maximum.
- ▶ p670 servers configured with Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418) can support up to 16 LPARs maximum.
- ▶ LPAR allocation, monitoring, and control is provided by the Hardware Management Console.
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of one processor is required per LPAR.
- ▶ A minimum of 1 GB of system memory is required per LPAR. While it is not required, 4 GB of system memory is recommended per LPAR to ensure adequate memory is available.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR on an individual slot basis.
- ▶ Integrated Ultra3 SCSI controllers located in the 7040-61D I/O drawers may be individually allocated to an LPAR. These integrated SCSI adapters each support one 4-pack disk backplane.
- ▶ While it is not mandatory, consideration should be given to allocating one half of a 7040-61D I/O drawer for each LPAR. This would provide one 10-slot PCI or PCI-X planar, two integrated SCSI controllers, and two 4-pack SCSI disk backplanes to the LPAR. This will help to ensure balanced I/O bandwidth for the LPAR configurations.

I/O drawers

This section provides information about the limitations of the pSeries 670 I/O drawers.

- ▶ The pSeries 670 utilizes the 7040-61D I/O drawer for directly attached PCI or PCI-X adapters and SCSI disk capabilities. The p670 requires a minimum of one 7040-61D drawer; a maximum of three I/O drawers can be attached.
- ▶ Each 7040-61D I/O drawer is divided into two separate halves. Each half contains ten blind-swap PCI or PCI-X slots and two Ultra3 SCSI 4-pack back planes for a total of twenty PCI slots and sixteen hot-swap disk bays per drawer.
- ▶ The first 7040-61D I/O drawer ordered with a complete new build pSeries 670 server order may be configured with only one I/O planar (FC 6563 or FC 6571) and two SCSI 4-pack backplanes.
- ▶ An I/O drawer configured with a single PCI planar (FC 6563) is connected to the p670 CEC using two 2-meter RIO cables (FC 3149). The 0.5 meter RIO cable (FC 3145) is not required for this configuration.
- ▶ An I/O drawer configured with a single PCI-X planar (FC 6571) is connected to the p670 CEC using two 1.75-meter RIO-2 cables (FC 3156).
- ▶ A single planar I/O drawer may be updated in the future through an MES. This update must include one additional I/O planar (FC 6563 or FC 6571) and two additional SCSI 4-pack backplanes (FC 6564).
- ▶ All other 7040-61D I/O drawers must be configured with two I/O planars and four SCSI 4-pack backplanes.
- ▶ p670 servers configured with a 7040-61D containing only one I/O planar (FC 6563 or FC 6571) are limited to a maximum of two media devices that must be mounted in the front half of the media drawer (FC 8692).
- ▶ One single-wide, blind-swap cassette (equivalent to those in FC 4599) is provided in each PCI or PCI-X slot of the 7040-61D. Cassettes not containing an adapter will be shipped with a “dummy” card installed to ensure proper environmental characteristics for the drawer. If additional single-wide, blind-swap cassettes are needed, FC 4599 should be ordered.
- ▶ All ten PCI or PCI-X slots on each I/O drawer planar are capable of supporting either 64-bit or 32-bit PCI or PCI-X adapters. Each FC 6563 planar provides seven 3.3V, 66 MHz, 64-bit PCI slots and three 5V, 32 MHz, 64-bit PCI slots. Each FC 6571 planar provides ten PCI-X slots capable of supporting 3.3V signaling PCI or PCI-X adapters operating at speeds up to 133 MHz.
- ▶ While it is not optimized, FC 6563 PCI and FC 6571 PCI-X planars may reside in the same 61D I/O drawer.

- ▶ Each I/O drawer planar incorporates two integrated Ultra3 SCSI adapters for direct attachment of the two 4-pack hot swap backplanes in that half of the drawer. These adapters do not support external SCSI device attachments.
- ▶ Each half of the 7040-61D I/O drawer is powered separately.
- ▶ I/O drawers occupy specific locations in the 7040-61R rack. Drawer placement is specified by the mounting location of the bottom of the drawer. Drawer placement is as follows:
 - Drawer #1 - Placement indicator FC 4609 - 9U position of 7040-61R rack
 - Drawer #2 - Placement indicator FC 4605 - 5U position of 7040-61R rack
 - Drawer #3 - Placement indicator FC 4601 - 1U position of 7040-61R rack

I/O drawer attachment

This section provides information about the limitations of the pSeries 670 I/O drawer attachment.

- ▶ FC 7040-61D I/O drawers are always connected to the p670 CEC by RIO (Remote I/O) or RIO-2 loops. 61D connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. p670 systems with non-looped configurations could experience degraded performance and serviceability.
- ▶ RIO loop connections operate at 500 MHz. RIO loops connect to the 7040-671 CEC using a Support Processor with Remote I/O Loop Attachment (FC 6404) (2 Loops) or Remote I/O Loop Adapter (FC 6410) (4 Loops).
- ▶ RIO-2 loop connections operate at 1 GHz. RIO-2 loops connect to the 7040-671 CEC 8 - Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418) (2 Loops) or Remote I/O-2 (RIO-2) Loop Adapter (FC 6419) (4 Loops).
- ▶ RIO loop attachment adapters (FC 6404 and FC 6410) are mutually exclusive with RIO-2 loop attachment adapters (FC 6418 and 6419).
- ▶ Up to three 7040-61D I/O drawers can be attached to the p670 depending on the server and attachment configuration. The Support Processor and Remote I/O Loop Attachment card (RIO FC 6404 or RIO-2 FC 6418) provides system control function as well as two RIO or RIO-2 loops for drawer attachment. The Remote I/O Loop Adapter (RIO FC 6410 or RIO-2 FC 6419) provides up to four additional RIO or RIO-2 loops.
- ▶ When utilizing a RIO attachment, the following number of I/O drawers are allowed per MCM:
 - One MCM allows attachment of up to two I/O drawers
 - Two MCMs allow attachment of up to three I/O drawers

- ▶ When utilizing RIO-2 attachment, a maximum of four loops are available for an I/O drawer attachment. Table 6-6 indicates the number of available I/O loops for each combination of MCMs vs. loop attachment adapters available.

Table 6-6 Number of available I/O loops for each combination of MCM

Number of MCM	With only FC 6418	With one FC 6419
One	2	N/A
Two	2	4

- ▶ When utilizing a RIO-2 loop attachment, the number of drawers that can be attached to a 7040-671 server is as follows:
 - Up to three I/O drawers, if all are attached using single-loop mode.
 - Up to two I/O drawers, if all are attached using dual-loop mode.
 - Up to three I/O drawers, if two are attached in single-loop mode and one is attached in dual-loop mode.
- ▶ 7040-61D I/O drawers may be connected to the 7040-671 CEC in either single-loop or dual-loop mode. Dual-loop mode is recommended whenever possible, as it provides the maximum bandwidth between the I/O drawer and the CEC.
- ▶ Single-loop mode connects an entire I/O drawer using one RIO or RIO-2 loop (2 ports) on the CEC. The two I/O planars in the I/O drawer are connected together using a short RIO or RIO-2 cable. Single-loop connection requires one loop (2 ports) per I/O drawer.
 - Single-loop mode connection is required for I/O drawers containing two FC 6563 PCI Planars whether connected using RIO or RIO-2 ports on the Model 671 CEC.
 - Single-loop mode is allowed for I/O drawers containing two FC 6571 PCI-X planars. However, dual-loop mode is recommended for this configuration.
- ▶ Dual-loop mode connects each I/O planar in the drawer to the CEC separately. Dual-loop connection requires RIO-2 connection and is not available with RIO connection. Each I/O planar is connected to the CEC using a separate RIO-2 loops. Dual-loop connection requires two loops (4 ports) per I/O drawer.
 - Dual-loop mode is required for I/O drawers containing one FC 6563 PCI I/O planar and one FC 6571 PCI-X I/O planar.
 - Dual-loop mode is recommended for all I/O drawers containing two FC 6571 I/O planars.

- ▶ To aid in manufacturing, it is recommended that initial p670 server orders containing both single and dual-looped drawers specify the single-looped drawers as the lowest numbered drawers in the configuration.
- ▶ FC 6563 PCI planars provide a RIO type cable connector for system attachment. When connecting to a CEC equipped with RIO ports, 2 meter RIO cables (FC 3149) are used. When connecting to a CEC equipped with RIO-2 ports, 3 meter RIO to RIO-2 cables (FC 3165) are used. Single-loop connection of two FC 6563 planars within an I/O drawer utilizes a 0.5 meter RIO cable (FC 3145) between the planars to complete the I/O loop.
- ▶ FC 6571 PCI-X I/O planars must be attached to the 7040-671 CEC using RIO-2 connections. RIO-2 connections for I/O drawers residing in the 7040-61R primary system rack are made using a 1.75 meter RIO-2 cables (FC 3156). Single-loop connection of two FC 6571 planars within an I/O drawer utilize a 1.2 meter RIO-2 cable (FC 3146) between the planars to complete the I/O loop

Disks and boot devices

This section provides information about the limitations of the pSeries 670 disks and boot devices.

- ▶ Minimum of two internal SCSI hard disks are required per p670 server. It is recommended that these disks be utilized as mirrored boot devices. These disks should be mounted in the first 7040-61D I/O drawer. This configuration provides service personnel the maximum amount of diagnostic information if the system encounters errors in the boot sequence.
- ▶ Boot support is also available from local SCSI, SSA, and Fibre Channel adapters, or from networks using ENET or token-ring adapters. The pSeries 670 does not support booting from FDDI adapters FC 2741 or FC 2742 located in 7040-61D I/O drawers.
- ▶ Consideration should also be given to the placement of AIX rootvg volume group in the first I/O drawer. This allows AIX to boot any time other I/O drawers are found offline during boot.
- ▶ If the boot source other than internal disk is configured, the supporting adapter should also be in the first I/O drawer.
- ▶ The p670 incorporates an Early Power Off Warning (EPOW) capability that assists in performing an orderly system shutdown in the event of a sudden power loss. IBM recommends use of the integrated battery backup features or an uninterruptible power system (UPS) to help ensure against loss of data due to power failures.

Rack and power

This section provides information about the limitations of the pSeries 670 rack and power.

- ▶ The 7040-61R rack is a 24-in. rack with an integrated power subsystem to support the pSeries 670 system. It provides 42 U of rack space.
- ▶ The 7040-61R rack can be utilized in multi-rack clusters. Rack indicators (such as FC 4651) will be used to identify the racks within the cluster and specify the desired rack in which to place various rack mounted components. When assigning rack indicators, the 7040-61R and its expansion rack FC 8691 (not available for p670 implementations) are considered a single rack.
- ▶ All 7040-61R racks and expansion feature racks must have door assemblies installed. The following door assemblies are available:
 - A sculptured black front door with copper accent (FC 6070) is required for the primary Model 61R rack.
 - A sculptured black front door (FC 6071) is required for the expansion rack feature, if installed.
 - An acoustic rear door (FC 6075) is available for the primary rack or the expansion feature rack. This door should be utilized for installations requiring a quieter environment.
 - A slim line rack door (FC 6074) is available for the primary or expansion rack feature. This door should be utilized for installations where system footprint is the primary consideration.
- ▶ If additional external communication and storage devices are required, 7014-T00 or T42 racks should be ordered. There is no limit on the quantity of 7014 racks allowed.
- ▶ The height of the 7040-61R rack (42 U) or expansion feature rack may require special handling when shipping by air or when moving under a low doorway.
- ▶ The 7040-61R rack always incorporates two bulk power assemblies for redundancy (FC 8690). These provide 350V DC power for devices located in the 61R rack and associated expansion rack feature. These bulk power assemblies are mounted in front and rear positions and occupy the top 8 U of the rack. To provide optimum system availability, these bulk power assemblies should be powered from separate power sources using separate line cords.
- ▶ Bulk power regulators (FC 6186) interface to the bulk power assemblies to ensure proper power is supplied to the systems components. Bulk power regulators are always installed in pairs in the front and rear bulk power assemblies to provide redundancy. The number of bulk power regulators required is configuration dependent based on the number of processor MCMs and I/O drawers installed.

Table 6-7 provides the number of bulk power regulators required for servers containing POWER4 1.1 GHz or 1.3 GHz processor MCMs.

Table 6-7 Bulk power regulators required with POWER4 1.1 or 1.3 GHz

Number of 7040-61D I/O drawer	One Power4 MCM	Two Power4 MCM
One	2	2
Two	2	2
Three	-	4

Table 6-8 provides the number of bulk power regulators required for servers containing POWER4+ 1.5 GHz MCMs.

Table 6-8 Bulk power regulators required for svrs w/ POWER4+ 1.5 GHz MCM

Number of 7040-61D I/O drawer	One Power4+ MCM	Two Power4+MCM
One	2	2
Two	2	4
Three	-	4

- ▶ Two Power Controller features (FC 6187) are required for all pSeries 670 systems, one per bulk power assembly. Each Power Controller provides power connections to support the system's four cooling fans. It also provides three additional connectors that can be utilized to attach DC power converters contained in the CEC. Ten additional connector locations are provided by the Power Distribution Assembly (FC 6188). Two Power Distribution Assemblies (one per Bulk Power Assembly) are required on each system. Two additional Power Distribution Assemblies may be added to provide more connections for large system configurations.
- ▶ Optional Integrated Battery Backup is available, if desired. The battery backup features protect against power line disturbances and provide sufficient power to allow an orderly system shutdown in the event that the power sources fail. The battery backup features each require 2 U of space in the primary 7040-61R system rack or in the expansion rack (FC 8691).

PCI and PCI-X slots and adapters

This section provides information about the limitations of the pSeries 670 PCI/PCI-X slots and adapters.

- ▶ System maximum limits for adapters and devices may not provide optimal system performance. These limits are given for connectivity and functionality assurance.
- ▶ Configuration limitations have been established to help ensure appropriate PCI or PCI-X bus loading, adapter addressing, and system and adapter functional characteristics when ordering I/O drawers. These I/O drawer limitations are in addition to individual adapter limitations shown in the feature descriptions section.
- ▶ Most PCI and PCI-X adapters for the pSeries 670 system are capable of being hot-plugged. Any PCI adapter supporting a boot device or system console should not be hot-plugged. The SP Switch Attachment Adapter (FC 8396) cannot be hot plugged.
- ▶ When configured with I/O drawers containing PCI planars (FC 6563), the combined quantities of the following adapters are limited to 30 adapters per pSeries 670 server, 10 per 7040-61D I/O drawer, and five per PCI planar (FC 6563). When configured with I/O drawers containing PCI-X planars (FC 6571), the combined quantities of these adapters are limited to 60 adapters per pSeries 670 server, 20 per 7040-61D I/O drawer, and 10 per PCI planar (FC 6571).
 - Turboways 622 Mbps PCI MMF ATM Adapter (FC 2946)
 - Gigabit Ethernet - SX PCI Adapter (FC 2969)
 - 10/100/1000 Base-T Ethernet PCI Adapter (FC 2975)
 - IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700)
 - IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701)
 - IBM 2-Port 10/100/1000 Base-TX E-net PCI-X Adapter (FC 5706)
 - IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707)
 - PCI Dual Channel Ultra3 SCSI Adapter (FC 6203)
 - Gigabit Fibre Channel Adapter for 64-bit PCI Bus (FC 6228)
 - Advanced SerialRAID Plus Adapter (FC 6230)
 - SP Switch PCI Attachment Adapter (FC 8396) (Note)
 - SP Switch2 PCI Attachment Adapter (FC 8397) (Note)
 - SP Switch2 PCI-X Attachment Adapter (FC 8398) (Note)

Note: Each instance of an SP Switch (FC 8396), SP Switch2 (FC 8397), or SP Switch2 PCI-X (FC 8398) adapter should be counted as two adapters when calculating the 30 or 60 adapters per p670 limitation.

- ▶ A combined maximum of 32 total 8-port and 128-port Asynchronous adapters (FC 2943 plus FC 2944) are allowed per 7040-671 server.
- ▶ A combined maximum of eight total SysKonnnect SK-Net FDDI adapters (FC 2741 plus FC 2742) are allowed per 7040-671 server.
- ▶ A combined maximum of eight total Serial HIPPI adapters (FC 2732 plus FC 2733) are allowed per 7040-671 server.
- ▶ A combined maximum of eight total POWER GXT135P Graphics Accelerators (FC 2848 plus FC 2849) are allowed per 7040-671 server.
- ▶ A combined total of 30 ATM adapters (FC 2946 plus FC 4953 plus FC 4957) are allowed per 7040-671 server.
- ▶ The SP Switch Attachment Adapter (FC 8396) and the SP Switch2 Attachment Adapters (FC 8397 and FC 8398) are mutually exclusive on a pSeries 670 system.
- ▶ A combined maximum of 14 total ESCON (FC 2751) and 4-port Multiprotocol Adapters (FC 2947) are allowed per LPAR.
- ▶ A maximum of eight total 4-port Multiprotocol Adapters (FC 2947) are allowed per LPAR when using SDLC, X.25, or bi-sync protocols.
- ▶ The maximum number of a specific PCI or PCI-X adapter allowed per p670 server may be less than the number allowed per I/O drawer multiplied by the maximum number of I/O drawers supported. The maximum “per drawer” adapter information is contained in the 7040-61D I/O drawer sales manual. Table 6-9 provides the maximum number of each adapter supported per p670 server and per LPAR within a p670 server.

Table 6-9 Supported adapters per LPAR within a p670

Feature code	Adapter	Max per LPAR	Max with only FC 6563 PCI planar & RIO	Max with only FC 6571 PCI-X planar & RIO 2
2732	IBM Short-wave Serial HIPPI	2	8	8
2733	IBM Long-wave Serial HIPPI	2	8	8
2737	Keyboard/Mouse Attachment Card	1	8	8
2741	SysKonnnect SK-NET FDDI-LP SAS	8	8	8
2742	SysKonnnect SK-NET FDDI-LP DAS	8	8	8
2848	POWER GXT135P Graphics Accelerator	1	8	8
2849	POWER GXT135P Graphics Accelerator W/Digital	1	8	8

Feature code	Adapter	Max per LPAR	Max with only FC 6563 PCI planar & RIO	Max with only FC 6571 PCI -X planar & RIO 2
2751	S/390 ESCON Channel Adapter	4	8	0 ²
2943	8-Port Asynchronous Adapter	16	16 ¹	16 ¹
2944	128-Port Asynchronous Controller	32	32 ¹	32 ¹
2946	Turboways 622 Mbps PCI MMF ATM	30	30	30
2947	IBM ARTIC960Hx 4-Port Multiprotocol	14	16	16
2962	2-Port Multiprotocol Adapter	18	18	18
2969	Gigabit Ethernet - SX PCI Adapter	30	30	30
2975	10/100/1000 Base-T Ethernet PCI Adapter	30	30	30
4953	IBM 64-Bit/66 MHz ATM 155 UTP Adapter	40	40	40
4957	IBM 64-Bit/66 MHz ATM 155 MMF Adapter	40	40	40
4959	IBM Token-Ring PCI Adapter	40	40	40
4960	IBM e-business Cryptographic Accelerator	24	24	24
4961	IBM 4-Port 10/100 Ethernet Adapter	20	20	40
4962	10/100 Mbps Ethernet PCI Adapter II	60	60	60
4963	PCI Cryptographic Coprocessor (FIPS-4)	4	24	24
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	30	30	30
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	30	30	30
5706	IBM 2-Port 10/100/1000 Base-TX E-net PCI-X Adapter	40	20	40
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	40	20	40
6203	PCI Dual Channel Ultra3 SCSI Adapter	30	30	30
6204	PCI Universal Differential Ultra SCSI Adapter	20	20	20

Feature code	Adapter	Max per LPAR	Max with only FC 6563 PCI planar & RIO	Max with only FC 6571 PCI -X planar & RIO 2
6206	PCI Single-Ended Ultra SCSI Adapter	2	2	0 ²
6228	Gigabit Fibre Channel Adapter	30	30	30
6230	Advanced SerialRAID Plus Adapter	30	30	30
8396	IBM SP Switch Attachment Adapter	1	4	0 ²
8397	IBM SP Switch2 Attachment Adapter	2	8	0 ²
8398	SP Switch2 PCI-X Attachment Adapter	2	0 ³	8

Notes:

1. Does not include FC 2943 and FC 2944 adapters configured as part of a FC 7315 or FC 7316 IBM Hardware Management Console for pSeries.
2. Requires FC 6563 PCI Planar.
3. Requires FC 6571 PCI-X Planar.

Clustered server environment

This section provides information about the limitations of the pSeries 670 in a clustered server environment.

- ▶ IBM @server pSeries and RS/6000 servers may be incorporated into a clustered environment operating under the control of the IBM Parallel Systems Support Programs for AIX (PSSP). Each cluster may contain a maximum of 16 pSeries 670 (7040-671) servers.
- ▶ The pSeries 670 can be divided into multiple logical partitions (LPARs). Up to 16 LPARs on a p670 can be configured to function in the clustered server environment. Each LPAR configured as a clustered server is considered as one of the servers that can be incorporated into the overall cluster. A maximum of 48 servers in a cluster may be created using pSeries server LPARs.
- ▶ Clustered server environments always incorporate a control workstation for cluster control. Clustered server environments are available in either switch-attached or non-switch attached implementations.
- ▶ SP Switch2 attached servers utilize one or two SP Switch2 Attachment adapters (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398) for high speed interconnection to the SP Switch2 mounted in a 9076 frame. This environment does not require a 9076 SP Node in the configuration.

- ▶ SP Switch attached servers utilize the SP Switch Attachment adapter (FC 8396) for high speed interconnection to the SP Switch mounted in a 9076 frame.
- ▶ The SP Switch (FC 8396), and Switch2 (FC 8397) adapters are “double wide” PCI cards and take up the space of two PCI slots. They must be mounted in the optional PCI Blind-Swap Cassette Kit for Double Wide Adapters (FC 4598).
- ▶ A non-switch attached environment utilizes an Ethernet connection to pass data between the clustered servers. This environment does not require a switch, 9076 Frame, or 9076 SP node in the configuration.
- ▶ The SP Switch2 supports two configurations, allowing communications over either one or two switch planes. Two switch plane configurations must incorporate two SP Switch2 Attachment Adapters (FC 8397) or Switch2 PCI-X Attachment Adapters (FC 8398) in each attached server. Implementations utilizing a one switch plane configuration are limited to a maximum of one Switch2 Attachment Adapter (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398) per attached server.
- ▶ Each pSeries 670 clustered server interfaces to the control workstation using Ethernet twisted pair connections. An Ethernet connection to the control workstations is required for each clustered server (or LPAR operating as a clustered server) as well as on the HMC controlling the p670 system. The 10/100 Mbps Ethernet PCI Adapter II (FC 4962) should be utilized for this connection for the p670 server. The Ethernet adapters supporting the LPARs and clustered servers are allowed in slots 8, 9, 18, and 19 of the 7040-61D I/O drawer when utilizing the SP Switch2 Attachment Adapter (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398). It can be installed in slots 1 through 7 or slots 11 through 17 with the SP Switch Attachment Adapter (FC 8396). The HMC incorporates an Ethernet connection in its basic configuration, which is utilized for this connection.
- ▶ An appropriate length switch cable is required to attach each switch attached server. This cable is ordered as a feature of the 9076 SP.
- ▶ Four SP Switch Attachment Adapters (FC 8396) are allowed per 7040-671 server. These adapters are located in the 7040-61D I/O drawers. A maximum of one adapter is allowed per FC 6563 I/O planar and two adapters per 7040-61D I/O drawer.
- ▶ Eight SP Switch2 Attachment Adapters (FC 8397) or Switch2 PCI-X Attachment Adapters (FC 8398) are allowed per 7040-671 server. These adapters are located in the 7040-61D I/O drawers. A maximum of two FC 8397 adapters is allowed per FC 6563 PCI I/O planar. A maximum of two FC 8398 adapters is allowed per FC 6571 PCI-X I/O planar.

- ▶ FC 8397 adapters are not allowed with FC 6571 PCI-X planars. FC 8398 adapters are not allowed with FC 6563 PCI planars.
- ▶ Some I/O adapters supported on the 7040-671 servers, when utilized in a non-clustered environment, are not supported in the clustered server environments and must be removed. Refer to the IBM RS/6000 9076-550 Sales Manual pages for a list of the currently supported adapters. For more information, refer to *RS/6000 SP Planning, Volume 1, Hardware and Physical Environment*, GA22-7280.
- ▶ Graphics adapters and natively attached displays are not supported in the clustered server environment. An ASCII terminal may be attached to servers in these environments but is not required. System control functions are provided by the SP control workstation.
- ▶ For additional information regarding server clustering and control workstations, visit the IBM PSSP Web site:
http://www.rs6000.ibm.com/resource/aix_resource/sp_books/
- ▶ The establishment of a trusted network between the control workstation (CWS) and the Hardware Management Console (HMC) is needed to bypass a security vulnerability in Cluster 1600 systems with pSeries 670 or 670 servers. Refer to the *PSSP Version 3.4 Read This First* document at the Web site:
http://www.ibm.com/servers/eserver/pseries/library/sp_books/pssp.html

IBM Capacity Upgrade on Demand for processors

This section provides information about the limitations of the pSeries 670 Capacity Upgrade on Demand for processors.

- ▶ Capacity Upgrade on Demand (CUoD) for processors is available for pSeries 670 (7040-671) servers. CUoD for processors allows inactive processors to be installed in the p670 server, which can be permanently activated by the customer as required.
- ▶ A single MCM is required on any 7040-671 server implementing CUoD for processors. Each CUoD MCM must have a minimum of four of the eight processors active.
- ▶ The first MCM in a 7040-671 servers implementing CUoD for POWER4 1.1 GHz processors must always have eight fully active processors. The second MCM may be implemented with CUoD for processors, if desired.
- ▶ 7040-671 servers utilizing POWER4+ 1.5 GHz processors may implement CUoD for processors on all installed MCMs if desired.
- ▶ Additional processors on the CUoD MCMs are activated in increments of two by ordering the appropriate activation feature number. If more than two

processors are to be activated at the same time, the activation code should be ordered multiple times.

- ▶ Upon receipt of an order for a CUoD for processors activation code, IBM will provide the customer a 32-character encrypted key. This key is entered into the system to activate the desired number of additional processors.
- ▶ CUoD processors that have not been activated are available to the 7040-671 server for dynamic processor sparing when running AIX 5L Version 5.2. If the server detects the impending failure of an active processor, it will attempt to activate one of the unused CUoD processors and add it to the system configuration. This helps to keep the server's processing power at full strength until a repair action can be scheduled.
- ▶ The following are required to support CUoD for processors:
 - AIX 5L Version 5.1 (with AIX 51-003 or higher Recommended Maintenance Package and PTFs U483632 and U483633) or later
 - Hardware Management Console for pSeries with Version 1.3 or later software
 - p670/670 (7040-671/681) system microcode (machine code) at the 11/14/2002 (RH021114) update or later
- ▶ AIX 5L Version 5.2 or later is recommended, as it provides enhanced CUoD function.
- ▶ CUoD on servers configured with AIX 5L Version 5.1 are subject to the following limitations:
 - In SMP mode, activation of standby processors requires a system IPL to enable additional processors.
 - AIX 5L Version 5.1 LPAR mode activation of standby processors requires a partition IPL to enable additional processors.
 - In mixed AIX 5L Version 5.1 and Version 5.2 LPAR mode, activation of standby processors for any partition running AIX 5L Version 5.1 requires a partition IPL. AIX 5L Version 5.2 partitions can dynamically allocate additional processors.
- ▶ Dynamic processor sparing is limited to systems and/or partitions that are running AIX 5L Version 5.2. Systems running only AIX 5L Version 5.1 provide processor sparing for a failed processor on a system IPL.

IBM On/Off Capacity on Demand for processors

This section provides information about the limitations of the pSeries 670 on/off capacity upgrade on demand for processors.

- ▶ On/Off Capacity on Demand (On/Off CoD) for processors is available for pSeries 670 (7040-671) servers. On/Off CoD for processors allows customers

to temporarily activate installed CUoD processors and later deactivate the resources as desired.

- ▶ Customers with installed but inactive CUoD processor resources can order On/Off CoD activation features. Each On/Off CoD activation ordered authorizes activation of two processors for 30 days of usage. These activations may be utilized for 30 consecutive days or turned on and off over a longer period of time if desired.
- ▶ Customers may order multiple On/Off CoD activation codes if desired. These may be utilized to activate multiple two processor increments for the 30 day period or to activate two processors for periods longer than thirty days. The HMC controlling the p670 server will be utilized to set the parameters desired when utilizing multiple On/Off CoD activation codes.

Trial Capacity on Demand (Trial CoD)

Trial Capacity on Demand (Trial CoD) is a function delivered with all pSeries servers supporting CUoD resources beginning May 30, 2003. Those servers with standby CUoD processors or memory will be capable of using a one-time, no-cost activation for a maximum period of 30 consecutive days. This enhancement allows for benchmarking of CUoD resources or can be used to provide immediate access to standby resources when the purchase of a permanent activation is pending.

Trial CoD is a complementary service offered by IBM. Although IBM intends to continue it for the foreseeable future, IBM reserves the right to withdraw Trial CoD at any time with or without notice.

IBM Capacity Upgrade on Demand for memory

This section provides information about the limitations of the pSeries 670 Capacity Upgrade on Demand for memory.

- ▶ Capacity Upgrade on Demand (CUoD) for memory is available for pSeries 670 (7040-671) servers. CUoD for processors allows inactive memory to be installed in the p670 server, which can be permanently activated by the customer as required.
- ▶ CUoD for memory utilizes inward facing 567 MHz memory cards in the following capacities:
 - 16 GB total card capacity with 8 GB active
 - 32 GB total card capacity with 16 GB active
- ▶ CUoD for memory cards must always be ordered in identical pairs.
- ▶ Memory configuration rules for the 7040-671 server apply to CUoD for memory cards as well as conventional memory cards. The memory

configuration rules are applied based upon the maximum capacity of the memory card.

- Apply 16 GB configuration rules for 16 GB CUoD for memory cards with less than 16 GB of active memory.
- Apply 32 GB configuration rules for 32 GB CUoD for memory cards with less than 32 GB of active memory.
- ▶ CUoD for memory may be utilized in any available memory position.
- ▶ Additional memory CUoD memory cards are activated in increments of 4 GB by ordering the appropriate activation feature number. If more than one 4 GB memory increment is to be activated at the same time, the activation code should be ordered multiple times.
- ▶ Upon receipt of an order for a CUoD for memory activation code, IBM will provide the customer a 32-character encrypted key. This key is entered into the system to activate the desired number of additional 4 GB memory increments.
- ▶ The following are required to support CUoD for memory:
 - AIX 5L Version 5.2 with 5200-01 Recommended Maintenance package (APAR IY39795) or later.
 - AIX 5L Version 5.1 for full system images only. A system reboot is required.
 - Hardware Management Console for pSeries with Version 1.3.2 or later software.
 - p670/670 (7040-671/681) system microcode (machine code) at the 05/2003 level update or later.

Dynamic logical partitioning

This section provides information about the limitations of the pSeries 670 dynamic logical partitioning.

- ▶ Dynamic logical partitioning (dynamic LPAR) allows system resources to be easily and quickly configured across multiple logical partitions (LPARs) on the p670 server.
- ▶ Dynamic LPAR can be used to add newly activated Capacity Upgrade on Demand processors and memory into the p670 configuration without requiring a system reboot.
- ▶ The following are required for dynamic LPAR:
 - AIX 5L Version 5.2 or later
 - Hardware Management Console for pSeries with Version 1.3 or later software

- p670/670 (7040-671/681) system microcode (machine code) at the 11/14/2002 (RH021114) update or later

Physical specifications/power and operating requirements

For physical specifications, power, and operating requirements, please refer to Appendix A, “Site and hardware planning information” on page 747.

6.4 7040-681 IBM @server pSeries 690

This section discusses the pSeries 690. Figure 6-19 show the pSeries 690.



Figure 6-19 Rack with expansion frame pSeries 690

6.4.1 Highlights pSeries 690

This sections explain the highlights of the pSeries 690 system.

- ▶ IBM SOI/copper chip technology
- ▶ Mainframe-inspired LPAR and self-managing capabilities
- ▶ Cost-efficient growth path to the future
- ▶ DDR System memory expandable from 8 GB to 512 GB

- ▶ Up to eight 7040-61D I/O drawers per server; each I/O drawer supports 20 blind-swap PCI or PCI-X bus slots and up to 16 Ultra3 SCSI disks that can be hot-plugged
- ▶ Hardware split into as many as sixteen LPARs, each functioning as a “computer within a computer” with its instance of the operating system
- ▶ Service processor, up to 160 hot-plug PCI and PCI-X slots, hot-plug redundant power supplies and cooling fans, and Dynamic Processor Deallocation
- ▶ Up to 32 dynamic partitions
- ▶ On demand capabilities including processor on demand, memory on demand, and On/Off Capacity on Demand
- ▶ AIX cluster or SP-attachment enabled
- ▶ AIX 5L and Linux
- ▶ 64-bit system scalability:
 - 8-way, 16-way, 24-way, or 32-way packaged on four Multi-Chip Modules (four or eight POWER4 or POWER4+ processors per MCM)
 - 1.1 GHz or 1.3 GHz POWER4 processors
 - 1.5 GHz or 1.7 GHz POWER4+ processors
 - 128 MB of L3 ECC cache per MCM

6.4.2 7040 pSeries 690 minimum and standard features

This section describes the system minimum, standard features of a pSeries 690 (see Table 6-10).

Table 6-10 The pSeries 670 minimum and standard feature

Specification	Detail
Processor	8-way SMP (one 8-way MCM); 1.5GHz or 1.7GHz POWER4+ or 1.1GHz or 1.3GHz POWER4
HPC minimum processor configuration	8-way SMP (two 4-way MCMs); 1.3GHz POWER4 HPC; 256 MB of L3 cache
L3 cache	128 MB
Memory (minimum)	8 GB
Internal disk drives	Two 36.4 GB Ultra SCSI
Disk bays	16 hot-swappable using one 7040-61D I/O drawer

Specification	Detail
Media bays	Five (four available)
Expansion slots	20 PCI or PCI-X (64-bit) using one 7040-61D I/O drawer
PCI bus width	32- and 64-bit
I/O adapters	Two integrated Ultra3 SCSI controllers
Standard Ports	Two serial ports for connecting Hardware Management Console for pSeries
Operating systems	AIX 5L Versions 5.1/5.2 SuSE Linux Enterprise Server Version 8 (runs in LPAR only)
Software Requirements	AIX 5L for POWER Version 5.1 with the 5100-02 Recommended Maintenance package (APAR IY 28102) or later, or AIX 5L Version 5.2 or later.

6.4.3 7040 pSeries 690 system expansion

Table 6-11 shows the possible maximum processor, memory, and storage configurations.

Table 6-11 Processor, memory, and storage configurations

System expansion	
SMP configuration	16-, 24-, 32-way SMP (two, three, or four 8-way MCMs); 1.5GHz or 1.7GHz POWER4+ or 1.1GHz or 1.3GHz POWER4
L3 cache	128 MB per MCM (256 MB maximum)
Memory	Up to 512 GB (ECC, Chipkill)
PCI expansion slots	Up to 160 adapters using two additional 7040-61D I/O drawers
Disk bay expansion	Up to 128 hot-swappable disk bays using two additional 7040-61D I/O drawers; up to 7.0 TB (36.4 GB, 73.4 GB, and 146.8 GB disk drives available)
Optional battery backup	Up to two
Attachment	SP System Attachment Adapter for use in a Cluster 1600 configuration

6.4.4 Configuration notes

In the following section, we discuss some configuration notes about the pSeries 690 system.

pSeries 690 Central Electronics Complex (CEC)

This section provides information about the limitations of the pSeries 690 CEC.

- ▶ The pSeries 690 CEC (7040-681) is a 17 U tall, 24-in., rack-mounted device. It houses the system processors, memory, system support processor, I/O drawer connection capability, and associated components. It is installed with the bottom of the drawer at placement indicator (FC 4418), the 18 U position of the 7040-61R rack.
- ▶ The p690 is powered by one to four multi-chip processor modules. Each Multi-Chip Module (MCM) contains either four or eight processors. The p690 is available in the following configurations:
 - 8-and 16-way 1.3 GHz, POWER4, HPC option processors (four processors per MCM)
 - 8-, 16-, 24-, 32-way 1.1 GHz, POWER4 processors (eight processors per MCM)
 - 8-, 16-, 24-, 32-way 1.3 GHz, POWER4, Turbo option processors (eight processors per MCM)
 - 8-, 16-, 24-, 32-way 1.5 GHz, POWER4+ processors (eight processors per MCM)
 - 8-, 16-, 24-, 32-way 1.7 GHz, POWER4+ Turbo option processors (eight processors per MCM)
- ▶ All processors in a p690 must operate at the same speed.
- ▶ Each processor contains 32 KB of data cache and 64 KB of instruction cache.
- ▶ Each processor MCM contains by 5.6 MB of Level 2 cache, which is shared between the processors. This provides 1.44 MB of L2 cache per processor for four processor MCMs and 0.72 MB of L2 cache per processor for eight processor MCMs.
- ▶ Each installed processor MCM is supported with 128 MB of Level 3 cache. This provides 32 MB of L3 cache per processor for four processor MCMs and 16 MB of L3 cache per processor for eight processor MCMs.
 - 1.1 GHz processors use L3 cache FC 4138
 - 1.3 GHz processors use L3 cache FC 4139
 - 1.5 GHz processors use L3 cache FC 4198
 - 1.7 GHz processors use L3 cache FC 4199

- ▶ A programmable processor clock card (FC 5251) is required to provide system and processor clocking.
- ▶ p690 servers with 1.5 GHz or 1.7 GHz POWER4+ processors must be equipped with feature Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418).
- ▶ p690 systems configured with two or three processor MCMs must have the empty processor positions populated with Processor Bus Pass Through Modules (FC 5257). No Processor Bus Pass Through Modules are required on systems with one processor MCM.
- ▶ p690 processor MCMs require capacitor books to operate. One capacitor book (FC 6198) is required for each two populated processor MCM positions.
- ▶ The p690 utilizes redundant power throughout its design. It implements redundant bulk power assemblies, bulk power regulators, power controllers, power distribution assemblies, DC power converters, and associated cabling.
- ▶ Power for the p690 CEC is supplied from DC bulk power assemblies in the 7040-61R rack. The bulk power is converted to the power levels required for the CEC using DC power converters. Two DC power converters (FC 6170) and attachment cables (FC 6181 and FC 6161) are required for the CEC and the first processor MCM, memory, and miscellaneous CEC components. Each additional MCM (second, third, or fourth) requires an additional DC power converter (FC 6189) and the appropriate power cable group (FC 6182, 6183, or 6184).
- ▶ One additional DC power converter (FC 6189) is required for any 32-way p690 server configured with four POWER4+ Turbo 1.7 GHz MCMs (FC 5246 and/or FC 7406). This DC power converter is attached to the bulk power assembly using power cable group FC 6202.
- ▶ One Media Drawer (FC 8692) is required for each p690 server. The media drawer is a 1U rack drawer, which is located directly below the CEC at the 17U location of the 7040-61R rack. It incorporates the system operator panel, a diskette drive, and up to four media devices.
- ▶ The media drawer is split into two separate sections. The front section houses the system operator panel, a diskette drive, and two media bays. Utilization of the rear section is optional (to provide space for two additional media devices, if desired).
- ▶ The front and rear sections of the media drawer must be connected to separate ports on SCSI PCI adapters in the first 7040-61D I/O drawer, which is located in drawer position FC 4609. When connecting to a SCSI adapter with a 68 pin P-style connector, the I/O Drawer to Media Drawer SCSI cable FC 2122 is used. When connecting to a SCSI adapter with a Mini-68 pin, the VHDCI connector converter cable FC 2118 is used in conjunction with

FC 2122. The SCSI adapters and cables for this attachment are ordered with the first 7040-61D I/O drawer.

- ▶ Power for each half of the media drawer is provided by the first I/O drawer using one or two I/O Drawer to Media Drawer Power Cables (FC 6179) ordered with the 7040-61D I/O drawer.
- ▶ One media device capable of reading CD-ROM media is required for each p690 system.
- ▶ The 4 mm tape drive (FC 6158) is allowed in the rear bays of the media drawer only if the operating environment is maintained at 24° C (75.2° F) or below and only on systems installed at 2,134 meters (7,000 ft.) altitude or below.
- ▶ The DVD-RAM drive (FC 2623), CD-ROM drive (FC 2624), and DVD-ROM drive (FC 2634) are limited to systems without Primary Integrated Battery Backup (FC 6200) or Redundant Integrated Battery Backup (FC 6201) when installed in the Media Drawer (FC 8692) rear bay locations.
- ▶ The IBM 80/160 GB Internal Tape Drive with VXA Technology (FC 6120) is limited to systems without Primary Integrated Battery Backup (FC 6200) or Redundant Integrated Battery Backup (FC 6201) installed, Media Drawer (FC 8692) front bay locations only, and a maximum system ambient operating temperature of 28° C (82.4° F) at a maximum operating altitude of 2,134 m (7,000 ft.). Lower altitudes have higher maximum ambient operating temperatures. Refer to the *7040 System Planning Guide* for additional details.
- ▶ USB keyboards supporting various language groups can be ordered with the 7040-61D I/O Drawer for use with native-attached displays on the pSeries 690. A three-button USB mouse connects to the USB keyboard.

pSeries 690 memory

This section provides information about the limitations of the pSeries 690 memory.

- ▶ The pSeries 690 has eight memory slots. Four memory slots utilize inward facing memory cards and four utilize outward facing memory cards. The inward facing memory slots are utilized by the first and second processor module positions while the outward facing memory slots support the third and fourth processor module positions.
- ▶ The number of active memory positions on a p690 system is related to the number of populated MCM locations. Two memory positions are activated for each MCM installed.
 - The first MCM installed activates the first and second inward facing memory positions.

- The second MCM installed activates the third and fourth inward facing memory positions.
- The third MCM installed activates the first and second outward facing memory positions.
- The fourth MCM installed activates the third and fourth outward facing memory positions.
- ▶ Minimum system memory is 8 GB. Maximum system memory is 512 GB.
- ▶ Memory is available in 4 GB, 8 GB, 16 GB, 32 GB, and 64 GB inward facing and outward facing increments.
- ▶ Memory is available in 500 MHz and 567 MHz speeds.
 - 500 MHz and 567 MHz memory can be simultaneously installed in a 7040-681 server. In this configuration, all memory will perform at 500 MHz.
 - 500 MHz memory allows maximum memory subsystem performance for 7040-681 servers with POWER4 or POWER4+ processor MCMs operating at speeds up to 1.5 GHz.
 - 567 MHz memory must be installed on servers utilizing POWER4+ 1.7 GHz processor MCMs to obtain maximum memory subsystem performance. Utilizing 500 MHz memory with these 1.7 GHz processors will result in degraded memory subsystem performance.
- ▶ 4 GB and 8 GB, 567 MHz memory features (FC 4480, FC 4481, FC 4482, and FC 4483) are available only on 7040-681 servers utilizing POWER4+ 1.7 GHz processor MCMs.
- ▶ The following memory configuration considerations are recommended to ensure the best performance is obtained for the pSeries 690 system:
 - Memory should be installed in identical pairs.
 - All activated memory locations should be populated with memory cards.
 - Memory sizes should be balanced as closely as possible across all populated MCM locations.

Table 6-12 on page 273 provides supported memory card placement for various memory requirements. The memory configurations with (*) are valid, but not recommended. You have to review application performance implications before ordering this memory configuration.

Table 6-12 Supported memory cards configurations

Number of MCMs	Memory size in GB	Memory slots 0 & 1	Memory slots 2 & 3	Memory slots 4 & 5	Memory slots 6 & 7
1	8	4 + 4			
	8 (*)	8 + 0			
	16	8 + 8			
	32	16 + 16			
	64	32 + 32			
	128	64+64			
2	8	4 + 4	0 + 0		
	16	4 + 4	4 + 4		
	16 (*)	8 + 8	0 + 0		
	24	8 + 8	4 + 4		
	32	8 + 8	8 + 8		
	48	16 + 16	8 + 8		
	64	16 + 16	16 + 16		
	96	32 + 32	16 + 16		
	128	32 + 32	32 + 32		
	192	64 + 64	32 + 32		
	256	64 + 64	64 + 64		

Number of MCMs	Memory size in GB	Memory slots 0 & 1	Memory slots 2 & 3	Memory slots 4 & 5	Memory slots 6 & 7
3	8	4 + 4	0 + 0	0 + 0	
	16	4 + 4	4 + 4	0 + 0	
	16 (*)	8 + 8	0 + 0	0 + 0	
	24	4 + 4	4 + 4	4 + 4	
	24 (*)	8 + 8	4 + 4	0 + 0	
	32	8 + 8	4 + 4	4 + 4	
	32 (*)	8 + 8	8 + 8	0 + 0	
	40	8 + 8	8 + 8	4 + 4	
	48	8 + 8	8 + 8	8 + 8	
	64	16 + 16	8 + 8	8 + 8	
	80	16 + 16	16 + 16	8 + 8	
	96	16 + 16	16 + 16	16 + 16	
	128	32 + 32	16 + 16	16 + 16	
	160	32 + 32	32 + 32	16 + 16	
	192	32 + 32	32 + 32	32 + 32	
	256	64 + 64	32 + 32	32 + 32	
320	64 + 64	64 + 64	32 + 32		
384	64 + 64	64 + 64	64 + 64		

Number of MCMs	Memory size in GB	Memory slots 0 & 1	Memory slots 2 & 3	Memory slots 4 & 5	Memory slots 6 & 7
4	8	4 + 4	0 + 0	0 + 0	0 + 0
	16	4 + 4	4 + 4	0 + 0	0 + 0
	16 (*)	8 + 8	0 + 0	0 + 0	0 + 0
	24	4 + 4	4 + 4	4 + 4	0 + 0
	24 (*)	8 + 8	4 + 4	0 + 0	0 + 0
	32	4 + 4	4 + 4	4 + 4	4 + 4
	32 (*)	8 + 8	4 + 4	4 + 4	0 + 0
	32 (*)	8 + 8	8 + 8	0 + 0	0 + 0
	40	8 + 8	4 + 4	4 + 4	4 + 4
	40 (*)	8 + 8	8 + 8	4 + 4	0 + 0
	48	8 + 8	8 + 8	4 + 4	4 + 4
	48 (*)	8 + 8	8 + 8	8 + 8	0 + 0
	56	8 + 8	8 + 8	8 + 8	4 + 4
	64	8 + 8	8 + 8	8 + 8	8 + 8
	80	16 + 16	8 + 8	8 + 8	8 + 8
	96	16 + 16	16 + 16	8 + 8	8 + 8
	112	16 + 16	16 + 16	16 + 16	8 + 8
	128	16 + 16	16 + 16	16 + 16	16 + 16
	160	32 + 32	16 + 16	16 + 16	16 + 16
	192	32 + 32	32 + 32	16 + 16	16 + 16
224	32 + 32	32 + 32	32 + 32	16 + 16	
256	32 + 32	32 + 32	32 + 32	32 + 32	
320	64 + 64	32 + 32	32 + 32	32 + 32	
384	64 + 64	64 + 64	32 + 32	32 + 32	
448	64 + 64	64 + 64	64 + 64	32 + 32	
512	64 + 64	64 + 64	64 + 64	64 + 64	

In the configurations listed previously in Table 6-12 on page 273, you notice that the first memory slots are populated with the largest memory boards in the first memory slots. This corresponds to the recommended order for inserting memory boards, starting with the largest boards in the first slots, and then installing the following boards in decreasing memory size order. IBM manufacturing will deliver the initial configuration of pSeries 690 with memory installed according to this rule. However, this order is only a recommendation, and there may be cases where it is not respected. For example, if a customer has a system initially configured with 48 GB memory as follows:

16 + 16	8 + 8		
---------	-------	--	--

then the customer orders an MES to increase the memory to 96 GB with one pair of 16 GB memory board and one pair of 8 GB memory board (assuming the system is populated with 4 MCM), and the memory configuration will become:

16 + 16	8 + 8	16 + 16	8 + 8
---------	-------	---------	-------

This is a valid and supported configuration. It would not be possible to reorder the memory boards, since the 8 GB memory boards initially installed in slots 2 and 3 are inward facing and could not be moved to slots 4, 5, 6, or 7 when the MES memory boards are installed.

Note: Consideration for system upgrades. When initially ordering a pSeries 690 configuration with the intent to order additional hardware later, you must pay attention to the supported memory configurations available with different numbers of MCM. For example, if you order a two MCM pSeries 690 with 256 GB of memory and you later want to order two extra MCM, you will have to also order at least 128 GB extra memory.

For more information on physical memory configuration and performance considerations, please refer the *IBM @server pSeries 690 Configuring for Performance* white paper, found at:

http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_config.htm
1

IBM Hardware Management Console for pSeries (HMC)

This section provides information about the limitations of the pSeries 690 HMC.

- ▶ The following Hardware Management Consoles can be used to control a p690 server:
 - HMC FC 7315 and FC 7316 purchased with a 7040-671 or 7040-681 server (withdrawn from marketing)
 - IBM 7315-C01 Hardware Management Console (withdrawn from marketing)
 - IBM 7315-C02 Hardware Management Console
 - IBM 7315-CR2 rack-mounted Hardware Management Console
- ▶ The HMC also provides a service focal point for the systems it controls. It is connected to the support (service) processor of the system by a dedicated serial link. The HMC provides tools for problem determination and service support, such as call-home and error log notification through an analog phone line.
- ▶ The HMC is a dedicated function device. It is utilized only for the control and service functions of the pSeries servers it serves. It is not available for use as a general purpose computing resource.
- ▶ Each HMC must be configured with a graphics display and a keyboard. The 7315-CR2 rack-mounted HMC can be connected to the IBM 7316-TF2 rack-mounted display and associated VGA switch. Several HMCs can, therefore, share the same display, mouse, and keyboard. Various native language keyboards are available for the HMC. However, the HMC displays menus and control information in English regardless of the keyboard utilized.
- ▶ The number of servers an HMC can control is dependent on the particular HMC and the server environment.
- ▶ HMC FC 7315 (withdrawn from marketing) can control up to four pSeries servers in a non-clustered environment.
- ▶ The 7315-C02 or 7315-C01 or FC 7316 HMC (ordered with a 7040-671 or 681 server) can control the following numbers of servers in a non-clustered environment:
 - 12 pSeries 670 and/or p690 servers with up to 64 LPARS or
 - 16 pSeries 655 servers with up to 32 LPARS or
 - 16 pSeries 630 or pSeries 650 servers with up to 64 LPARS or
 - 16 pSeries servers and 64 LPARS in a mixed server environment. A mixed server environment can contain a combined maximum of eight p670 and/or p690 servers.

- ▶ If redundant HMC function is desired, the pSeries 670 can be attached to two separate hardware management consoles.
- ▶ An Ethernet connection between the HMC and each active partition on the pSeries server is recommended for all installations and required for dynamic LPAR configurations. This connection is utilized to provide:
 - Additional systems management, such as Web-based System Manager of AIX in the individual partitions.
 - Collection and passing of hardware service events to the HMC for automatic notification of error conditions to IBM.
 - Total system inventory collection.
- ▶ The HMC provides two integrated serial ports. One serial port is required for each pSeries server controlled by the HMC. An additional serial port is required for modem attachment if the Service Agent call-home function is implemented. When more than two serial ports are required, 8-port (FC 2943) and 128-port (FC 2944) asynchronous adapters should be utilized.
- ▶ 128-Port Asynchronous Adapters are not supported on the FC 7315 HMC.
- ▶ Up to two 8-port or 128-port asynchronous adapters are allowed per HMC.
- ▶ Each p690 CEC must be connected to a Hardware Management Console (HMC) for system control, LPAR (logical partitioning), Capacity Upgrade on Demand, and service functions. The HMC is capable of supporting multiple p690 systems.
- ▶ The p690 provides two special asynchronous ports specially dedicated for HMC attachment. The connectors on these ports are differentiated from standard asynchronous ports by providing a 9-pin female connector vs. the standard 9-pin male connector. The connection between the HMC and the p690 CEC is through cables available in 6-meter (FC 8120) and 15-meter (FC 8121) lengths.
- ▶ The p690 also provides two asynchronous ports, which are available for attachment of ASCII terminals and devices. These two asynchronous ports are under the control of the system's service processor and are not available for functions such as HACMP heartbeat cabling or UPS control, which require fully dedicated ports. The service processor menus may be accessed through an ASCII terminal attached to one of the asynchronous ports or by an internal firmware controlled "wrap" connection to the HMC. When not in use for accessing service processor menus, the asynchronous ports are available for general purpose uses, which do not require dedicated ports.

Logical partitioning (LPAR)

This section provides information about the limitations of the pSeries 690 logical partitioning.

- ▶ The pSeries 690 can be divided into as many as 32 logical partitions. System resources can be dedicated to each LPAR.
- ▶ Servers configured with Support Processor with Remote I/O Loop Attachment (FC 6404) can support up to 16 LPARs maximum.
- ▶ Servers configured with Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (6418) can support up to 32 LPARs maximum.
- ▶ LPAR allocation, monitoring, and control is provided by the Hardware Management Console.
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of one processor and 1 GB of memory is required per LPAR.
- ▶ A minimum of 4 GB of system memory is recommended per LPAR.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR on an individual slot basis.
- ▶ Integrated Ultra3 SCSI controllers located in the 7040-61D I/O drawers may be individually allocated to an LPAR. These integrated SCSI adapters each support one 4-pack disk backplane.
- ▶ While it is not mandatory, consideration should be given to allocating one half of a 7040-61D I/O drawer for each LPAR. This would provide one 10-slot PCI or PCI-X planar, two integrated SCSI controllers, and two 4-Pack SCSI disk backplanes to the LPAR. This will help to ensure balanced I/O bandwidth for the LPAR configurations.

I/O drawers

This section provides information about the limitations of the pSeries 690 I/O drawers.

- ▶ The pSeries 690 utilizes the 7040-61D I/O drawer for directly attached PCI or PCI-X adapters and SCSI disk capabilities. The p690 requires a minimum of one 7040-61D drawer; a maximum of eight I/O drawers can be attached.
- ▶ Each 7040-61D I/O drawer is divided into two separate halves. Each half contains ten blind-swap PCI or PCI-X slots and two Ultra3 SCSI 4-pack back planes for a total of twenty PCI slots and sixteen hot-swap disk bays per drawer.
- ▶ One single-wide, blind-swap cassette (equivalent to those in FC 4599) is provided in each PCI or PCI-X slot of the 7040-61D. Cassettes not containing an adapter will be shipped with a “dummy” card installed to ensure proper

environmental characteristics for the drawer. If additional single-wide, blind-swap cassettes are needed, FC 4599 should be ordered.

- ▶ All ten PCI or PCI-X slots on each I/O drawer planar are capable of supporting either 64-bit or 32-bit PCI or PCI-X adapters. Each FC 6563 planar provides seven 3.3V, 66 MHz, 64-bit PCI slots and three 5V, 32 MHz, 64-bit PCI slots. Each FC 6571 planar provides ten PCI-X slots capable of supporting 3.3V signaling PCI or PCI-X adapters operating at speeds up to 133 MHz.
- ▶ While it is not optimized, FC 6563 PCI and FC 6571 PCI-X planars may reside in the same 61D I/O drawer.
- ▶ Each I/O drawer planar incorporates two integrated Ultra3 SCSI adapters for direct attachment of the two 4-pack hot swap backplanes in that half of the drawer. These adapters do not support external SCSI device attachments.
- ▶ Each half of the 7040-61D I/O drawer is powered separately.
- ▶ I/O drawers occupy specific locations in the 7040-61R rack or expansion feature rack. Drawer placement is specified by the mounting location of the bottom of the drawer. Drawer placement is as follows:
 - Drawer #1 - Placement indicator FC 4609 - 9 U position of 7040-61R rack
 - Drawer #2 - Placement indicator FC 4605 - 5 U position of 7040-61R rack
 - Drawer #3 - Placement indicator FC 4601 - 1 U position of 7040-61R rack
 - Drawer #4:
 - FC Placement indicator FC 4613 - 13 U position of 7040-61R rack, if Integrated Battery Backup feature (FC 6200) is not installed.
 - Placement indicator FC 4409 - 9U position of expansion rack, if Integrated Battery Backup feature (FC 6200) is installed at location FC 4613.
 - Drawer #5 - Placement indicator FC 4401 - 1 U position of expansion rack
 - Drawer #6 - Placement indicator FC 4405 - 5 U position of expansion rack
 - Drawer #7 - Placement indicator FC 4413 - 13 U position of the expansion rack
 - Drawer #8 - Placement indicator FC 4419 - 19 U position of the expansion rack

I/O drawer attachment

This section provides information about the limitations of the pSeries 690 I/O drawer attachment.

- ▶ FC 7040-61D I/O drawers are always connected to the p690 CEC using RIO (Remote I/O) or RIO-2 loops. 61D connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. p690 systems with non-looped configurations could experience degraded performance and serviceability.
- ▶ RIO loop connections operate at 500 MHz. RIO loops connect to the 7040-681 CEC by a Support Processor with Remote I/O Loop Attachment (FC 6404) (2 Loops) or Remote I/O Loop Adapter (FC 6410) (4 Loops).
- ▶ RIO-2 loop connections operate at 1 GHz. RIO-2 loops connect to the 7040-681 CEC Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418) (2 Loops), or Remote I/O-2 (RIO-2) Loop Adapter (FC 6419) (4 Loops).
- ▶ RIO Loop Attachment Adapters (FC 6404 and FC 6410) are mutually exclusive with RIO-2 Loop Attachment Adapters (FC 6418 and 6419).
- ▶ Up to eight 7040-61D I/O drawers can be attached to the p690, depending on the server and attachment configuration. The Support Processor and Remote I/O Loop Attachment card (RIO FC 6404 or RIO-2 FC 6418) provides system control function as well as two RIO or RIO-2 loops for drawer attachment. Each Remote I/O Loop Adapter (RIO FC 6410 or RIO-2 FC 6419) provides up to four additional four RIO or RIO-2 loops.
- ▶ When utilizing RIO attachment, the following number of I/O drawers are allowed per MCM:
 - One MCM allows attachment of up to two I/O drawers.
 - Two MCMs allow attachment of up to four I/O drawers.
 - Three MCMs allow attachment of up to six I/O drawers.
 - Four MCMs allow attachment of up to eight I/O drawers.
- ▶ When utilizing RIO-2 attachments, a maximum of 14 loops are available for I/O drawer attachment. Table 6-13 indicates the number of available I/O loops for each combination of MCMs vs. loop attachment adapters available.

Table 6-13 Available I/O loops

Number of MCM	With only 6418	With one 6419	With two 6419	With three 6419
One	2	N/A	N/A	N/A
Two	2	4	6	N/A

Number of MCM	With only 6418	With one 6419	With two 6419	With three 6419
Three	2	6	8	12
Four	2	6	10	14

- ▶ When utilizing a RIO-2 loop attachment, the number of drawers that can be attached to a 7040-681 server is as follows:
 - Up to eight I/O drawers, if all are attached by a single-loop mode.
 - Up to seven I/O drawers, if all are attached by a dual-loop mode.
 - Up to eight I/O drawers, with a combination of single-loop and dual-loop attachment. This requires that at least two drawers are connected in single-loop mode.
- ▶ 7040-61D I/O drawers may be connected to the 7040-681 CEC in either single-loop or dual-loop mode. Dual-loop mode is recommended whenever possible, as it provides the maximum bandwidth between the I/O drawer and the CEC.
- ▶ Single-loop mode connects an entire I/O drawer to the CEC using one RIO or RIO-2 loop (two ports). The two I/O planars in the I/O drawer are connected together using a short RIO or RIO-2 cable. Single-loop connection requires one loop (two ports) per I/O drawer.
 - Single-loop mode connection is required for I/O drawers containing two FC 6563 PCI planars whether connected by a RIO or RIO-2 ports on the Model 681 CEC.
 - Single-loop mode is allowed for I/O drawers containing two FC 6571 PCI-X planars. However dual-loop mode is recommended for this configuration.
- ▶ Dual-loop mode connects each I/O planar in the drawer to the CEC separately. Dual-loop connection requires RIO-2 connection and is not available with RIO connection. Each I/O planar is connected to the CEC using a separate RIO-2 loops. Dual-loop connection requires two loops (four ports) per I/O drawer.
 - Dual-loop mode is required for I/O drawers containing one FC 6563 PCI I/O planar and one FC 6571 PCI-X I/O planar.
 - Dual-loop mode is recommended for all I/O drawers containing two FC 6571 I/O planars.
- ▶ To aid in manufacturing, it is recommended that initial p690 server orders containing both single-and dual-looped drawers specify the single-looped drawers as the lowest numbered drawers in the configuration.
- ▶ PCI I/O planars (FC 6563) provide a RIO type cable connector for system attachment. When connecting to a CEC equipped with RIO ports, 2 meter

RIO cables (FC 3149) are used. When connecting to a CEC equipped with RIO-2 ports, 3 meter RIO to RIO-2 cables (FC 3165) are used. A single-loop connection of two FC 6563 planars within an I/O drawer utilizes a 0.5 meter RIO cable (FC 3145) between the planars to complete the I/O loop.

- ▶ PCI-X I/O planars (FC 6571) must be attached to the 7040-681 CEC using RIO-2 connections. RIO-2 connections for I/O drawers residing in the 7040-61R primary system rack are made using 1.75 meter RIO-2 cables (FC 3156). I/O drawers residing in the secondary rack feature connect to the CEC by a 2.5 meter RIO-2 cables (FC 3168). A single-loop connection of two FC 6571 planars within an I/O drawer utilize a 1.2 meter RIO-2 cable (FC 3146) between the planars to complete the I/O loop.

Disks and boot devices

This section provides information about the limitations of the pSeries 690 disks and boot devices.

- ▶ A minimum of two internal SCSI hard disks are required per p690 server. It is recommended that these disks be utilized as mirrored boot devices. These disks should be mounted in the first 7040-61D I/O drawer. This configuration provides service personnel the maximum amount of diagnostic information if the system encounters errors in the boot sequence.
- ▶ Boot support is also available from local SCSI, SSA, and Fibre Channel adapters, or from networks using ENET or token-ring adapters. The pSeries 690 does not support booting from FDDI adapters FC 2741 or FC 2742 located in 7040-61D I/O drawers.
- ▶ Consideration should also be given to the placement of the AIX rootvg volume group in the first I/O drawer. This allows AIX to boot any time other I/O drawers are found offline during boot.
- ▶ If the boot source other than internal disk is configured, the supporting adapter should also be in the first I/O drawer.
- ▶ The p690 incorporates an Early Power Off Warning (EPOW) capability that assists in performing an orderly system shutdown in the event of a sudden power loss. IBM recommends the use of the integrated battery backup features or an uninterruptible power system (UPS) to help ensure against loss of data due to power failures.

Rack and power

This section provides information about the limitations of the pSeries 690 rack and power.

- ▶ The 7040-61R rack is a 24-in. rack with an integrated power subsystem to support the pSeries 690 system. It provides 42 U of rack space.

- ▶ An expansion rack feature (FC 8691) is available if additional 24-in. rack space is required. To install the expansion rack feature, the side cover of the 7040-61R rack is removed, the expansion rack is connected to the 61R, and the side cover is placed on the exposed side of the feature rack. Power for components in the expansion rack is provided from the bulk power assemblies in the Model 61R rack.
- ▶ Over time, the 7040-61R rack will be utilized in multi-rack clusters. Rack indicators (such as FC 4651) will be used to identify the racks within the cluster and specify the desired rack in which to place various rack mounted components. When assigning rack indicators, the 7040-61R and its expansion rack FC 8691 are considered a single rack.
- ▶ All 7040-61R racks and expansion feature racks must have door assemblies installed. The following door assemblies are available:
 - A sculptured black front door with copper accent (FC 6070) is required for the primary Model 61R rack.
 - A sculptured black front door (FC 6071) is required for the expansion rack feature, if installed.
 - An acoustic rear door (FC 6075) is available for the primary rack or the expansion feature rack. This door should be utilized for installations requiring a quieter environment.
 - A slim line rack door (FC 6074) is available for the primary or expansion rack feature. This door should be utilized for installations where system footprint is the primary consideration.
- ▶ If additional external communication and storage devices are required, 7014-T00 or T42 racks should be ordered. There is no limit on the quantity of 7014 racks allowed.
- ▶ The height of the 7040-61R rack (42 U) or expansion feature rack may require special handling when shipping by air or when moving under a low doorway.
- ▶ The 7040-61R rack always incorporates two bulk power assemblies for redundancy (FC 8690). These provide 350V DC power for devices located in the 61R rack and associated expansion rack feature. These bulk power assemblies are mounted in the front and rear positions and occupy the top 8 U of the rack. To provide optimum system availability, these bulk power assemblies should be powered from separate power sources using separate line cords.
- ▶ Bulk power regulators (FC 6186) interface to the bulk power assemblies to ensure proper power is supplied to the systems components. Bulk power regulators are always installed in pairs in the front and rear bulk power assemblies to provide redundancy. The number of bulk power regulators required is configuration dependent, based on the number of processor MCMs and I/O drawers installed.

- ▶ Table 6-14 provides the number of bulk power regulators required for servers containing POWER4 1.1 GHz or 1.3 GHz processor MCMs.

Table 6-14 Bulk power regulators with POWER4 1.1 or 1.3GHz

Number of 7040-61D I/O drawer	One Power4 MCM	Two Power4 MCM	Three Power4 MCM	Four Power4 MCM
One	2	2	4	4
Two	2	2	4	4
Three	-	4	4	4
Four	-	4	4	4
Five	-	-	6	6
Six	-	-	6	6
Seven	-	-	-	6
Eight	-	-	-	6

- ▶ Table 6-15 provides the number of bulk power regulators required for servers containing POWER4+ 1.5 GHz or 1.7 GHz processor MCMs.

Table 6-15 Bulk power regulators with POWER4+ 1.5 or 1.7GHz

Number of 7040-61D I/O drawer	One Power4 MCM	Two Power4 MCM	Three Power4 MCM	Four Power4 MCM
One	2	2	4	4
Two	2	2	4	4
Three	-	4	4	4
Four	-	4	4	6
Five	-	4	6	6
Six	-	-	6	6
Seven	-	-	-	6
Eight	-	-	-	6

- ▶ Two Power Controller features (FC 6187) are required for all pSeries 690 systems, one per bulk power assembly. Each Power Controller provides power connections to support the system's four cooling fans. It also provides three additional connectors that can be utilized to attach DC power converters

contained in the CEC. Ten additional connector locations are provided by the Power Distribution Assembly (FC 6188).

- ▶ Two Power Distribution Assemblies (one per Bulk Power Assembly) are required on each system.
- ▶ Additional Power Distribution Assemblies are added to provide more connections for large system configurations.
 - Four power distribution assemblies are required for configurations implementing one or more 7040-61D I/O drawers in the Expansion Rack (FC 8691).
 - Four power distribution assemblies are required for configurations containing four 7040-61D I/O drawers and four FC 6189 CEC DC Power Converters.
 - Six power distribution assemblies are required for configurations containing eight 7040-61D I/O drawers and four FC 6189 CEC DC Power Converters without battery backup features FC 6200 and FC 6201 present.
- ▶ Optional Integrated Battery Backup is available, if desired. The battery backup features protect against power line disturbances and provide sufficient power to allow an orderly system shutdown in the event that the power sources fail. The battery backup features each require 2 U of space in the primary 7040-61R system rack or in the expansion rack (FC 8691).

PCI and PCI-X slots and adapters

This section provides information about the limitations of the pSeries 690 PCI/PCI-X slots and adapters.

- ▶ System maximum limits for adapters and devices may not provide optimal system performance. These limits are given for connectivity and functionality assurance.
- ▶ Configuration limitations have been established to help ensure appropriate PCI or PCI-X bus loading, adapter addressing, and system and adapter functional characteristics when ordering I/O drawers. These I/O drawer limitations are in addition to individual adapter limitations shown in the feature descriptions section.
- ▶ Most PCI and PCI-X adapters for the pSeries 690 system are capable of being hot-plugged. Any PCI adapter supporting a boot device or system console should not be hot-plugged. The SP Switch Attachment Adapter (FC 8396) cannot be hot plugged.
- ▶ When configured with I/O drawers containing PCI planars (FC 6563), the combined quantities of the following adapters are limited to 80 adapters per pSeries 690 server, 10 per 7040-61D I/O drawer, and 5 per PCI planar

(FC 6563). When configured with I/O drawers containing PCI-X planars (FC 6571), the combined quantities of these adapters are limited to 140 adapters per pSeries 690 server, 20 per 7040-61D I/O drawer, and 10 per PCI planar (FC 6571).

- Turboways 622 Mbps PCI MMF ATM Adapter (FC 2946)
- Gigabit Ethernet - SX PCI Adapter (FC 2969)
- 10/100/1000 Base-T Ethernet PCI Adapter (FC 2975)
- IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700)
- IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701)
- IBM 2-Port 10/100/1000 Base-TX E-net PCI-X Adapter (FC 5706)
- IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707)
- PCI Dual Channel Ultra3 SCSI Adapter (FC 6203)
- Gigabit Fibre Channel Adapter for 64-bit PCI Bus (FC 6228)
- Advanced SerialRAID Plus Adapter (FC 6230)
- SP Switch PCI Attachment Adapter (FC 8396)(1)
- SP Switch2 PCI Attachment Adapter (FC 8397)(1)
- SP Switch2 PCI-X Attachment Adapter (FC 8398)(1)

Note: Each instance of an SP Switch (FC 8396), SP Switch2 (FC 8397), or SP Switch2 PCI-X (FC 8398) adapter should be counted as two adapters when calculating the 80 or 140 adapters per p690 limitation.

- ▶ A combined maximum of 32 total 8-Port and 128-Port Asynchronous adapters (FC 2943 plus FC 2944) are allowed per 7040-681 server.
- ▶ A combined maximum of eight total SysKonnnect SK-Net FDDI adapters (FC 2741 plus FC 2742) are allowed per 7040-681 server.
- ▶ A combined maximum of eight total Serial HIPPI adapters (FC 2732 plus FC 2733) are allowed per 7040-681 server.
- ▶ A combined maximum of eight total POWER GXT135P Graphics Accelerators (FC 2848 plus FC 2849) are allowed per 7040-681 server.
- ▶ A combined total of 40 ATM adapters (FC 2946 plus FC 4953 plus FC 4957) are allowed per 7040-681 server.
- ▶ The SP Switch Attachment Adapter (FC 8396) and the SP Switch2 Attachment Adapters (FC 8397 and FC 8398) are mutually exclusive on a pSeries 690 system.

- ▶ A combined maximum of 14 total ESCON (FC 2751) and 4-port Multiprotocol Adapters (FC 2947) are allowed per LPAR.
- ▶ A maximum of eight total 4-port Multiprotocol Adapters (FC 2947) are allowed per LPAR when using SDLC, X.25, or bi-sync protocols.
- ▶ The maximum number of a specific PCI or PCI-X adapter allowed per p690 server may be less than the number allowed per I/O drawer multiplied by the maximum number of I/O drawers supported. The maximum “per drawer” adapter information is contained in the 7040-61D I/O drawer sales manual. Table 6-16 provides the maximum number of each adapter supported per p690 server and per LPAR within a p690 server.

Table 6-16 Adapter supported per LPAR within a p690

Feature code	Adapter	Max per LPAR	Max with only FC 6563 PCI planar & RIO	Max with only FC 6571 PCI-X planar & RIO 2
2732	IBM Short-wave Serial HIPPI	2	8	8
2733	IBM Long-wave Serial HIPPI	2	8	8
2737	Keyboard/Mouse Attachment Card	1	8	8
2741	SysKonnnect SK-NET FDDI-LP SAS	8	8	8
2742	SysKonnnect SK-NET FDDI-LP DAS	8	8	8
2848	POWER GXT135P Graphics Accelerator	1	8	8
2849	POWER GXT135P Graphics Accelerator W/Digital	1	8	8
2751	S/390 ESCON Channel Adapter	4	8	0 ²
2943	8-Port Asynchronous Adapter	16	16 ¹	16 ¹
2944	128-Port Asynchronous Controller	32	32 ¹	32 ¹
2946	Turboways 622 Mbps PCI MMF ATM	40	40	40
2947	IBM ARTIC960Hx 4-Port Multiprotocol	14	16	16
2962	2-Port Multiprotocol Adapter	18	18	18
2969	Gigabit Ethernet - SX PCI Adapter	30	30	30
2975	10/100/1000 Base-T Ethernet PCI Adapter	30	30	30
4953	IBM 64-Bit/66MHz ATM 155 UTP Adapter	40	40	40

Feature code	Adapter	Max per LPAR	Max with only FC 6563 PCI planar & RIO	Max with only FC 6571 PCI -X planar & RIO 2
4957	IBM 64-Bit/66MHz ATM 155 MMF Adapter	40	40	40
4959	IBM Token-Ring PCI Adapter	40	40	40
4960	IBM e-business Cryptographic Accelerator	32	32	32
4961	IBM 4-Port 10/100 Ethernet Adapter	20	20	40
4962	10/100 Mbps Ethernet PCI Adapter II	80	80	140
4963	PCI Cryptographic Coprocessor (FIPS-4)	4	32	32
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	80	80	140
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	80	80	140
5706	IBM 2-Port 10/100/1000 Base-TX E-net PCI-X Adapter	40	40	80
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	40	40	40
6203	PCI Dual Channel Ultra3 SCSI Adapter	30	30	30
6204	PCI Universal Differential Ultra SCSI Adapter	20	20	20
6206	PCI Single-Ended Ultra SCSI Adapter	2	2	0 ²
6228	Gigabit Fibre Channel Adapter	80	80	80
6230	Advanced SerialRAID Plus Adapter	60	64	64
8396	IBM SP Switch Attachment Adapter	1	8	0 ²
8397	IBM SP Switch2 Attachment Adapter	2	32	0 ²
8398	SP Switch2 PCI-X Attachment Adapter	2	0 ³	32

Notes:

1. Does not include FC 2943 and FC 2944 adapters configured as part of a FC 7315 or FC 7316 IBM Hardware Management Console for pSeries.
2. Requires FC 6563 PCI Planar.
3. Requires FC 6571 PCI-X Planar.

Clustered server environment

This section provides information about the limitations of the pSeries 690 in a clustered server environment.

- ▶ IBM @server pSeries and RS/6000 servers may be incorporated into a clustered environment operating under the control of the IBM Parallel Systems Support Programs for AIX (PSSP). Each cluster may contain a maximum of 16 pSeries 670 (7040-671) and pSeries 690 (7040-681) servers.
- ▶ The pSeries 690 can be divided into multiple logical partitions (LPARs). Up to 16 LPARs on a p690 can be configured to function in the clustered server environment. Each LPAR configured as a clustered server is considered as one of the servers that can be incorporated into the overall cluster. A maximum of 48 servers in a cluster may be created using pSeries server LPARs.
- ▶ Clustered server environments always incorporate a control workstation for cluster control. Clustered server environments are available in either switch-attached or non-switch attached implementations.
- ▶ SP Switch2 attached servers utilize one or two SP Switch2 Attachment adapters (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398) for high speed interconnection to the SP Switch2 mounted in a 9076 frame. This environment does not require a 9076 SP Node in the configuration.
- ▶ SP Switch attached servers utilize the SP Switch Attachment adapter (FC 8396) for high speed interconnection to the SP Switch mounted in a 9076 frame.
- ▶ The SP Switch (FC 8396), and Switch2 (FC 8397) adapters are “double wide” PCI cards and take up the space of two PCI slots. They must be mounted in the optional PCI Blind-Swap Cassette Kit for Double Wide Adapters (FC 4598).
- ▶ A non-switch attached environment utilizes an Ethernet connection to pass data between the clustered servers. This environment does not require a switch, 9076 Frame, or 9076 SP node in the configuration.
- ▶ The SP Switch2 supports two configurations, allowing communications over either one or two switch planes. Two switch plane configurations must incorporate two SP Switch2 Attachment Adapters (FC 8397) or Switch2 PCI-X Attachment Adapters (FC 8398) in each attached server. Implementations utilizing a one switch plane configuration are limited to a maximum of one Switch2 Attachment Adapter (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398) per attached server.
- ▶ Each pSeries 690 clustered server interfaces to the control workstation using Ethernet twisted pair connections. An Ethernet connection to the control workstations is required for each clustered server (or LPAR operating as a clustered server) as well as on the HMC controlling the p690 system. The

10/100 Mbps Ethernet PCI Adapter II (FC 4962) should be utilized for this connection for the p690 server. The Ethernet adapters supporting the LPARs and clustered servers are allowed in slots 8, 9, 18, and 19 of the 7040-61D I/O drawer when utilizing the SP Switch2 Attachment Adapter (FC 8397) or Switch2 PCI-X Attachment Adapter (FC 8398). It can be installed in slots 1 through 7 or slots 11 through 17 with the SP Switch Attachment Adapter (FC 8396). The HMC incorporates an Ethernet connection in its basic configuration, which is utilized for this connection.

- ▶ An appropriate length switch cable is required to attach each switch attached server. This cable is ordered as a feature of the 9076 SP.
- ▶ Eight SP Switch Attachment Adapters (FC 8396) are allowed per 7040-681 server. These adapters are located in the 7040-61D I/O drawers. A maximum of one adapter is allowed per FC 6563 I/O planar and two adapters per 7040-61D I/O drawer.
- ▶ Thirty-two SP Switch2 Attachment Adapters (FC 8397) or Switch2 PCI-X Attachment Adapters (FC 8398) are allowed per 7040-681 server. These adapters are located in the 7040-61D I/O drawers. A maximum of two FC 8397 adapters is allowed per FC 6563 PCI I/O planar. A maximum of two FC 8598 adapters is allowed per FC 6571 PCI-X I/O planar.
- ▶ Some I/O adapters supported on the 7040-681 servers when utilized in a non-clustered environment are not supported in the clustered server environments and must be removed. Refer to the IBM RS/6000 9076-550 Sales Manual pages for a list of the currently supported adapters. For more information, refer to the document *RS/6000 SP Planning, Volume 1, Hardware and Physical Environment, GA22-7280*.
- ▶ FC 8397 adapters are not allowed with FC 6571 PCI-X planars. FC 8398 adapters are not allowed with FC 6563 PCI planars.
- ▶ Graphics adapters and natively attached displays are not supported in the clustered server environment. An ASCII terminal may be attached to servers in these environments but is not required. System control functions are provided by the SP control workstation.
- ▶ For additional information regarding server clustering and control workstations, visit the IBM PSSP Web site:
http://www.rs6000.ibm.com/resource/aix_resource/sp_books/
- ▶ The establishment of a trusted network between the control workstation (CWS) and the Hardware Management Console (HMC) is needed to bypass a security vulnerability in Cluster 1600 systems with pSeries 690 or 670 servers. Refer to the *PSSP Version 3.4 Read This First* document at the Web site:
http://www.ibm.com/servers/eserver/pseries/library/sp_books/pssp.html

Capacity Upgrade on Demand for processors

This section provides information about the limitations of the pSeries 690 capacity upgrade on demand for processors.

- ▶ Capacity Upgrade on Demand (CUoD) for processors is available for pSeries 690 (7040-681) servers. CUoD for processors allows inactive processors to be installed in the p690 server, which can be permanently activated by the customer as required.
- ▶ A minimum of eight active processors are required on any 7040-681 server implementing CUoD for processors. Each CUoD MCM must have a minimum of four of the eight processors active.
- ▶ The first MCM in a 7040-681 servers implementing CUoD for POWER4 1.1 GHz or 1.3 GHz processors must always have eight fully active processors. The second, third, and fourth MCMs may be implemented with CUoD for processors, if desired.
- ▶ Servers utilizing POWER4+ 1.5 GHz or 1.7 GHz processors may implement CUoD for processors on all installed MCMs, if desired.
- ▶ Additional processors on the CUoD MCMs are activated in increments of two by ordering the appropriate activation feature number. If more than two processors are to be activated at the same time, the activation code should be ordered multiple times.
- ▶ Upon receipt of an order for a CUoD for processors activation code, IBM will provide the customer with a 32-character encrypted key. This key is entered into the system to activate the desired number of additional processors.
- ▶ CUoD processors that have not been activated are available to the 7040-681 server for dynamic processor sparing when running AIX 5L Version 5.2. If the server detects the impending failure of an active processor, it will attempt to activate one of the unused CUoD processors and add it to the system configuration. This helps to keep the server's processing power at full strength until a repair action can be scheduled.
- ▶ The following are required to support CUoD for processors:
 - AIX 5L Version 5.1 (with AIX 51-003 or higher Recommended Maintenance Package and PTFs U483632 and U483633) or later
 - Hardware Management Console for pSeries with Version 1.3 or later software
 - p670/690 (7040-671/681) system microcode (machine code) at the 11/14/2002 (RH021114) update or later
- ▶ AIX 5L Version 5.2 or later is recommended as it provides enhanced CUoD function.

- ▶ CUoD on servers configured with AIX 5L Version 5.1 are subject to the following limitations:
 - In SMP mode, activation of standby processors requires a system IPL to enable additional processors.
 - AIX 5L Version 5.1 LPAR mode activation of standby processors requires a partition IPL to enable additional processors.
 - In mixed AIX 5L Version 5.1 and 5.2 LPAR mode, activation of standby processors for any partition running AIX 5L Version 5.1 requires a partition IPL. AIX 5L Version 5.2 partitions can dynamically allocate additional processors.
- ▶ Dynamic processor sparing is limited to systems and/or partitions that are running AIX 5L Version 5.2. Systems running only AIX 5L Version 5.1 provide processor sparing for a failed processor on a system IPL.

Server On/Off Capacity on Demand for processors

This section provides information about the limitations of the pSeries 690 on/off capacity upgrade on demand for processors.

- ▶ On/Off Capacity on Demand (On/Off CoD) for processors is available for pSeries 690 (7040-681) servers. On/Off CoD for processors allows customers to temporarily activate installed CUoD processors and later deactivate the resources as desired.
- ▶ Customers with installed but inactive CUoD processor resources can order On/Off CoD activation features. Each On/Off CoD activation ordered authorizes activation of two processors for 30 days of usage. These activations may be utilized for 30 consecutive days or turned on and off over a longer period of time if desired.
- ▶ Customers may order multiple On/Off CoD activation codes if desired. These may be utilized to activate multiple two processor increments for the thirty day period or to activate two processors for periods longer than thirty days. The HMC controlling the p690 server will be utilized to set the parameters desired when utilizing multiple On/Off CoD activation codes.

Trial Capacity on Demand (Trial CoD)

Trial Capacity on Demand (Trial CoD) is a function delivered with all pSeries servers supporting CUoD resources beginning May 30, 2003. Those servers with standby CUoD processors or memory will be capable of using a one-time, no-cost activation for a maximum period of 30 consecutive days. This enhancement allows for benchmarking of CUoD resources or can be used to provide immediate access to standby resources when the purchase of a permanent activation is pending.

Trial CoD is a complementary service offered by IBM. Although IBM intends to continue it for the foreseeable future, IBM reserves the right to withdraw Trial CoD at any time with or without notice.

Capacity Upgrade on Demand for memory

This section provides information about the limitations of the pSeries 690 capacity upgrade on demand for memory.

- ▶ Capacity Upgrade on Demand (CUoD) for memory is available for pSeries 690 (7040-681) servers. CUoD for processors allows inactive memory to be installed in the p690 server, which can be permanently activated by the customer as required.
- ▶ CUoD for memory utilizes inward facing and outward facing 567 MHz memory cards in the following capacities:
 - 16 GB total card capacity with 8 GB active
 - 32 GB total card capacity with 16 GB active
- ▶ CUoD for memory cards must always be ordered in identical pairs.
- ▶ Memory configuration rules for the 7040-681 server apply to CUoD for memory cards as well as conventional memory cards. The memory configuration rules are applied based upon the maximum capacity of the memory card.
 - Apply 16 GB configuration rules for 16 GB CUoD for memory cards with less than 16 GB of active memory.
 - Apply 32 GB configuration rules for 32 GB CUoD for memory cards with less than 32 GB of active memory.
- ▶ CUoD for memory may be utilized in any available memory position.
- ▶ Additional memory CUoD memory cards are activated in increments of 4 GB by ordering the appropriate activation feature number. If more than one 4 GB memory increment is to be activated at the same time, the activation code should be ordered multiple times.
- ▶ Upon receipt of an order for a CUoD for memory activation code, IBM will provide the customer with a 32-character encrypted key. This key is entered into the system to activate the desired number of additional 4 GB memory increments.
- ▶ The following are required to support CUoD for memory:
 - AIX 5L Version 5.2 with 5200-01 Recommended Maintenance package (APAR IY39795) or later
 - Hardware Management Console for pSeries with Version 1.3.2 or later software

- p670/690 (7040-671/681) system microcode (machine code) at the 05/2003 level update or later

Dynamic logical partitioning

This section provides information about the limitations of the pSeries 670 dynamic logical partitioning.

- ▶ Dynamic logical partitioning (dynamic LPAR) allows system resources to be easily and quickly configured across multiple logical partitions (LPARs) on the p690 server.
- ▶ Dynamic LPAR can be used to add newly activated Capacity Upgrade on Demand processors and memory into the p690 configuration without requiring a system reboot.
- ▶ The following are required for dynamic LPAR:
 - AIX 5L for POWER Version 5.2 or later
 - Hardware Management Console for pSeries with Version 1.3 or later software
 - p670/690 (7040-671/681) system microcode (machine code) at the 11/14/2002 (RH021114) update or later

Physical specifications/power and operating requirement

For physical specifications, power, and operating requirements, please refer to Appendix A, “Site and hardware planning information” on page 747.

6.5 Features of the 7040 pSeries 670/690

This section lists the internal features that can be added to the pSeries 670 and pSeries 690 configuration. The status of a feature is indicative of these qualifications:

- | | |
|----------|--|
| X | Available. Indicates features that are available and orderable on the specified model. |
| R | Must be removed. Indicates this feature is not supported and must be removed during a model conversion. |
| W | Withdrawn. Indicates a feature that is no longer available to order. Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, mice, language specify codes, and power cords, are not included. |

Table 6-17 shows the features of the:

- ▶ pSeries 670/690 I/O drawer 61D
- ▶ pSeries 670/690 Rack 61R
- ▶ pSeries 670 CEC 671
- ▶ pSeries 690 CEC 681

Table 6-17 Feature code description pSeries 670/690.

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
Host attachment adapters					
2751	S/390 ESCON Channel PCI Adapter	X			
8396	IBM RS/6000 SP System Attachment Adapter	X			
8397	SP Switch2 PCI Attachment Adapter	X			
8398	SP Switch2 PCI-X Attachment Adapter	X			
Asynchronous adapters					
2943	8-Port Asynchronous Adapter EIA-232/RS-422, PCI bus	X		X	X
2944	128-Port Asynchronous Controller, PCI bus	X		X	X
Asynchronous cables					
2936	Asynchronous Cable EIA-232/V.24	X			
2934	Asynchronous Terminal/Printer Cable EIA-232	X		X	X
3926	Asynchronous Printer/Terminal Cable, 9-pin to 25-pin, 4M	X			
8131	128-Port Asynchronous Controller Cable, 4.5 Meter	X		X	X
8132	128-Port Asynchronous Controller Cable, 23cm (9-in.)	X		X	X
8121	Attachment Cable, Hardware Management Console to Host, 15-Meter			X	X
8120	Attachment Cable, Hardware Management Console to Host, 6-Meter			X	X
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232	X		X	X
8133	RJ-45 to DB-25 Converter Cable	X		X	X
3925	Serial Port Converter Cable, 9-Pin to 25-Pin	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
SCSI adapters					
6204	PCI Universal Differential Ultra SCSI Adapter	X			
6206	PCI Single-Ended Ultra SCSI Adapter	X			
6203	PCI Dual Channel Ultra3 SCSI Adapter	X			
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter	X			
5712	PCI-X Dual Channel Ultra320 SCSI Adapter	X			
Serial adapters					
6231	128 MB DRAM Option Card	X			
6230	Advanced SerialRAID Plus Adapter	X			
6235	32 MB Fast-Write Cache Option Card	X			
Fibre Channel adapters					
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus	W			
6239	2 Gigabit Fibre Channel PCI-X Adapter	X			
Fibre cables					
2456	LC-SC 50 Micron Fiber Converter Cable	X			
2459	LC-SC 62.5 Micron Fiber Converter Cable	X			
Graphics adapters					
2848	POWER GXT135P Graphics Accelerator	X			
2849	POWER GXT135P Graphics Accelerator with Digital Support	X			
Display cable					
4242	6 Foot Extender Cable for Displays (15 pin D-shell to 15 pin D-shell)	X			
LAN adapters					
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	X			
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	X			
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	X			
4959	IBM Token-Ring PCI Adapter	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
2946	Turboways 622 Mbps PCI MMF ATM Adapter	X			
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	X			
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	X			
4962	10/100 Mbps Ethernet PCI Adapter II	X		X	X
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	W			
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	W			
2969	Gigabit Ethernet - SX PCI Adapter	X			
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	X			
2975	10/100/1000 Base-T Ethernet PCI Adapter	X			
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	X			
WAN adapters					
2962	2-Port Multiprotocol PCI Adapter	X			
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	X			
WAN cables					
2861	ARTIC960Hx 4-Port EIA-232 Cable	X			
2864	ARTIC960Hx 4-Port V.35 (DTE) Cable	X			
2863	ARTIC960Hx 4-Port X.21 Cable	X			
2951	Cable, V.24/EIA-232	X			
2952	Cable, V.35	X			
2953	Cable, V.36/EIA-499	X			
2954	Cable, X.21	X			
Adapter mounting hardware					
4599	PCI Blind Swap Cassette Kit, Single Wide Adapters, Universal	X			
4598	PCI Blind Swap Cassette Kit, Double Wide Adapters, Universal	X			
Miscellaneous					
4963	PCI Cryptographic Coprocessor (FIPS-4)	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
4960	IBM e-business Cryptographic Accelerator	X			
2733	IBM Long-wave Serial HIPPI PCI Adapter for RS/6000	W			
2732	IBM Short-wave Serial HIPPI PCI Adapter for RS/6000	X			
2737	Keyboard/Mouse Attachment Card - PCI	X			
Miscellaneous internal system					
6410	Remote I/O Loop Adapter, Four Loop			X	X
6419	Remote I/O-2 (RIO-2) Loop Adapter, Four Loop			X	X
6404	Support Processor with Remote I/O Loop Attachment, Two Loops			X	X
6418	Support Processor with Remote I/O-2 (RIO-2) Loop Attachment, Two Loops			X	X
3155	Diskette Drive Cable			X	X
3255	Operator Panel Attachment Cable			X	X
3125	Serial to Serial Port Cable for Rack/Rack	X			
IIIDE internal					
4253	SCSI-to-IDE Interface Bridge			X	X
SCSI external					
2424	0.6M 16-bit SCSI-2 System to System Cable	X			
2425	2.5M 16-bit SCSI-2 System to System Cable	X			
2116	PCI SCSI Adapter To 2-Port, 16-Bit Differential Ext. Device Cable	W			
2118	Converter Cable, VHDCI to P, Mini-68 pin to 68 pin, 0.3M	X			
2122	SCSI Cable - I/O Drawer to Media Drawer	X			
2114	PCI SCSI Adapter 16-Bit Differential External Y Cable	X			
Rack related					
6121	I/O Drawer Attachment Cable Group, Drawer Position FC 4609	X			
6122	I/O Drawer Attachment Cable Group, Drawer Position FC 4605	X			
6123	I/O Drawer Attachment Cable Group, Drawer Position FC 4601	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
6124	I/O Drawer Attachment Cable Group, Drawer Position FC 4613	X			
6125	I/O Drawer Attachment Cable Group, Drawer Position FC 4401	X			
6126	I/O Drawer Attachment Cable Group, Drawer Position FC 4405	X			
6127	I/O Drawer Attachment Cable Group, Drawer Position FC 4409	X			
6128	I/O Drawer Attachment Cable Group, Drawer Position FC 4413	X			
6129	I/O Drawer Attachment Cable Group, Drawer Position FC 4419	X			
7827	I/O Drawer Attachment Cable Group, Drawer Position FC 4613, Alternate	X			
7848	I/O Drawer Attachment Cable Group, Drawer Position FC 4419, Alternate	X			
3164	RIO to RIO-2 Remote I/O Cable, 1.0M	X			
3146	RIO-2 (Remote I/O-2) Cable, 1.2M	X			
3165	RIO to RIO-2 Remote I/O Cable, 3.0M	X			
3156	RIO-2 (Remote I/O-2) Cable, 1.75M	X			
3168	RIO-2 (Remote I/O-2) Cable, 2.5M	X			
3145	Remote I/O Cable, 0.5M	X			
3149	Remote I/O Cable, 2M	X			
SCSI disks					
3157	18.2 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W			
3158	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	X			
3277	36.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly	X			
3159	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	X			
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly	X			
3275	146.8 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly	X			
Monitors					
3636	L200p Flat Panel Monitor	X			
3635	T210 Flat-Panel Monitor	W			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
3628	IBM P260/P275 Color Monitor, Business Black, and Cable	X		X	X
3627	IBM P76/P77 Color Monitor, Business Black, Captured Cable	W		W	W
3637	IBM T541H 15" TFT Color Monitor, Business Black, Captured Cable	X		X	X
3630	IBM P260/P275 Color Monitor, Pearl White, and Cable				W
3629	IBM P76/P77 Color Monitor, Pearl White, Captured Cable				W
Consoles					
7315	Hardware Management Console I				W
7316	Hardware Management Console II, Black			W	W
Memory cards					
4196	4 GB Memory Card, Inward Facing			X	X
4197	4 GB Memory Card, Outward Facing				X
4480	4 GB Memory Card, 567MHz, Inward Facing				X
4481	4 GB Memory Card, 567MHz, Outward Facing				X
4181	8 GB Memory Card, Inward Facing			X	X
4182	8 GB Memory Card, Outward Facing				X
4482	8 GB Memory Card, 567MHz, Inward Facing				X
4483	8 GB Memory Card, 567MHz, Outward Facing				X
4183	16 GB Memory Card, Inward Facing			W	W
4184	16 GB Memory Card, Outward Facing				W
4484	16 GB Memory Card, 567MHz, Inward Facing			X	X
4485	16 GB Memory Card, 567MHz, Outward Facing				X
4188	32 GB Memory Card, Inward Facing			W	W
4189	32 GB Memory Card, Outward Facing				W
4486	32 GB Memory Card, 567MHz, Inward Facing			X	X
4487	32 GB Memory Card, 567MHz, Outward Facing				X
4488	64 GB Memory Card, 567MHz, Inward Facing			X	X

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
4489	64 GB Memory Card, 567MHz, Outward Facing				X
Memory CUoD					
7060	4 GB CUoD Memory Activation for FC 7050 and FC 7051			X	X
7061	4 GB CUoD Memory Activation for FC 7054 and FC 7055			X	X
7050	16 GB 567MHz CUoD Memory Card, Inward Facing, 8 GB Active			X	X
7051	16 GB 567MHz CUoD Memory Card, Outward Facing, 8 GB Active				X
7054	32 GB 567MHz CUoD Memory Card, Inward Facing, 16 GB Active			X	X
7055	32 GB 567MHz CUoD Memory Card, Outward Facing, 16 GB Active				X
Cache					
8011	128 MB Level 3 Cache, 400 MHz - Upgrade Indicator				X
8013	128 MB Level 3 Cache, 500 MHz - Upgrade Indicator				X
4138	128 MB Level 3 Cache (4 X 32 MB), 400 MHz				X
4139	128 MB Level 3 Cache (4 X 32 MB), 433 MHz				X
4198	128 MB Level 3 Cache (4 x 32 MB), 500 MHz				X
4199	128 MB Level 3 Cache (4 x 32 MB), 567MHz				X
4699	128 MB Level 3 Cache (4 x 32 MB), 567MHz - Alternate				X
Power					
6187	Power Controller, Four Cooling Fans, Three DC Power Converter Connections		X		
6188	Power Distribution Assembly		X		
6186	Bulk Power Regulator		X		
6200	Integrated Battery Backup, Primary		X		
6201	Integrated Battery Backup, Redundant		X		
CEC related					
6189	DC Power Converter, CEC, Additional			X	X
6198	Capacitor Book, Two Processor Modules			X	X

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
6170	Power Converter Assembly, Central Electronics Complex			X	X
Drawer related					
6161	Cable Group, Power Controller to CEC and Fans			X	X
6181	Power Cable Group, CEC to Power Controller, First Processor Module			X	X
6162	Interface Cable, Service Processor to Power Subsystem			X	X
6172	Power Converter Assembly, I/O Drawer	X			
6241	Cable, Integrated Battery to Bulk Power Regulator, Expansion Rack		X		
6240	Cable, Integrated Battery Backup to Bulk Power Regulator, Primary Rack		X		
6182	Power Cable Group, CEC to Power Controller, Second Processor Module			X	X
6202	Power Cable Group, CEC to Power Controller, Additional Power Option				X
6183	Power Cable Group, CEC to Power Controller, Third Processor Module				X
6184	Power Cable Group, CEC to Power Controller, Fourth Processor Module				X
6179	Power Cable, I/O Drawer to Media Drawer	X			
SCSI CD/DVD					
2624	32X (Max) SCSI-2 CD-ROM Drive			X	X
2623	4.7 GB SCSI-2 DVD-RAM Drive (Black Bezel)			X	X
SCSI tapes					
6158	20/40 GB 4mm Internal Tape Drive			X	X
6120	IBM 80/160 GB Internal Tape Drive with VXA Technology			X	X
IDE CD/DVD					
2634	16X/48X(max) IDE DVD-ROM Drive			X	X
Disk bays					
6564	Ultra3 SCSI 4-Pack Hot Swap Back Plane	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
I/O planars					
6563	I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports	X			
6571	I/O Drawer PCI-X Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports	X			
CEC planars					
9670	Manufacturing Optimization Indicator			X	R
8007	System CEC Backplane - Upgrade Indicator				X
8015	System CEC Backplane Replacement - Upgrade Indicator				X
6565	Backplane, Central Electronics Complex				X
Software preload					
5005	Software Preinstall			X	X
Processors					
8012	POWER4+ 1.5GHz MCM - Upgrade Indicator				X
8009	POWER4 1.1GHz First MCM - Upgrade Indicator				X
5242	8-Way POWER4 Processor (First)				X
5240	4-Way POWER4 HPC Processor Option (First)				X
5244	8-Way POWER4 Turbo Processor Option (First)				X
5245	8-way POWER4 Turbo Processor Option (First) MES Feature				X
8010	POWER4 1.1GHz Second MCM - Upgrade Indicator				X
5253	4-Way POWER4 Processor with 128 MB L3 Cache			X	R
5256	8-Way POWER4 Processor with 128 MB L3 Cache			X	R
5247	4-Way POWER4+ Processor with 128 MB L3 Cache			X	R
5243	8-Way POWER4+ Processor Option				X
5248	8-Way POWER4+ Processor with 128 MB L3 Cache			X	R
5246	8-Way POWER4+ Turbo Processor Option				X
5209	PowerPC 800 MHz Processor with 2 MB of L3 Cache				X
5252	8-Way POWER4 Processor				X

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
5250	4-Way POWER4 HPC Processor Option				X
5254	8-Way POWER4 Turbo Processor Option				X
8014	Processor Bus Pass Through Module - Upgrade Indicator				X
5257	Processor Bus Pass Through Module				X
Processor CUoD					
7410	Two Processor Activation for CUoD Processor FC 7400			X	R
7412	Two Processor Activation for CUoD Processor FC 7402			X	R
7413	Two Processor Activation for CUoD Processor FC 7403				X
7414	Two Processor Activation for CUoD Processor FC 7404				X
7415	Two Processor Activation for CUoD Processor FC 7405				X
7416	Two Processor Activation for CUoD Processor FC 7406				X
7420	On/Off CoD 30-Day Two Processor Activation for Processor FC 7400			X	R
7422	On/Off CoD 30-Day Two Processor Activation for Processor FC 7402			X	R
7423	On/Off CoD 30-Day Two Processor Activation for Processor FC 7403				X
7424	On/Off CoD 30-Day Two Processor Activation for Processor FC 7404				X
7425	On/Off CoD 30-Day Two Processor Activation for Processor FC 7405				X
7426	On/Off CoD 30-Day Two Processor Activation for Processor FC 7406				X
7402	8-Way POWER4 CUoD Processor with 128 MB L3 Cache, 4-Way Active			X	R
7400	8-Way POWER4+ CUoD Processor with 128 MB L3 Cache, 4-Way Active			X	R
7403	8-Way POWER4 CUoD Processor Option, 4-Way Active				X
7405	8-Way POWER4 Turbo CUoD Processor Option, 4-Way Active				X
7404	8-Way POWER4+ CUoD Processor 4-Way Active				X
7406	8-Way POWER4+ Turbo CUoD Processor 4-Way Active				X
Processor clock					
5251	Processor Clock Card, Programmable			X	X

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
Rack related					
8690	Bulk Power Assembly, Redundant		X		
8691	Expansion Rack, 24", 42U		X		
Rack drawer					
8692	Media Drawer, Operator Panel, Diskette			X	X
Rack miscellaneous hardware					
6070	Front Door, Black with Copper Accent, Primary Rack		X		
6071	Front Door, Sculptured Black, Secondary Rack		X		
6075	Rear Door, Acoustic, Primary or Secondary Rack		X		
6074	Rear Door, Slim Line, Primary or Secondary Rack		X		
9004	Southern Hemisphere Designator for Monitors	X		X	X
Rack content					
0191	Rack Content Specify: 7040/61D - 4U		X		
0209	Rack Content Specify: 7040/671 - 17U		X		
0190	Rack Content Specify: 7040/681 - 17U		X		
0193	Rack Content Specify: 7040/61R FC 8690 - 8U		X		
0192	Rack Content Specify: Media Drawer FC 8692 - 1U		X		
Rack ID					
4401	Drawer Placement Indicator, 1U Position, Expansion Rack	X			
4601	Drawer Placement Indicator, 1U Position, Primary Rack	X			
4405	Drawer Placement Indicator, 5U Position, Expansion Rack	X			
4605	Drawer Placement Indicator, 5U Position, Primary Rack	X			
4409	Drawer Placement Indicator, 9U Position, Expansion Rack	X			
4609	Drawer Placement Indicator, 9U Position, Primary Rack	X			
4413	Drawer Placement Indicator, 13U Position, Expansion Rack	X			
4613	Drawer Placement Indicator, 13U Position, Primary Rack	X			

Feature code	Description	6 1 D	6 1 R	6 7 1	6 8 1
4618	Drawer Placement Indicator, 18U Position, Primary Rack			X	X
4419	Drawer Placement Indicator, 19U Position, Expansion Rack	X			
4619	Drawer Placement Indicator, 19U Position, Primary Rack	X			
4623	Drawer Placement Indicator, 23U Position, Primary Rack	X			
4627	Drawer Placement Indicator, 27U Position, Primary Rack	X			
4631	Drawer Placement Indicator, 31U Position, Primary Rack	X			
4651	Rack Indicator, Rack # 1	X	X	X	X
4652	Rack Indicator, Rack # 2	X	X	X	X
4653	Rack Indicator, Rack # 3	X	X	X	X
4654	Rack Indicator, Rack #4	X	X	X	X
4655	Rack Indicator, Rack #5	X	X	X	X
4656	Rack Indicator, Rack #6	X	X	X	X
4657	Rack Indicator, Rack #7	X	X	X	X
4658	Rack Indicator, Rack #8	X	X	X	X
4659	Rack Indicator, Rack #9	X	X	X	X
4660	Rack Indicator, Rack #10	X	X	X	X
4661	Rack Indicator, Rack #11	X	X	X	X
4662	Rack Indicator, Rack #12	X	X	X	X
4663	Rack Indicator, Rack #13	X	X	X	X
4664	Rack Indicator, Rack #14	X	X	X	X
4665	Rack Indicator, Rack #15	X	X	X	X
4666	Rack Indicator, Rack #16	X	X	X	X



Racks and rack solution

This chapter contains information about the following racks: the Rack-Mounted Flat Panel Console and Uninterruptible Power System (UPS). The following topics are discussed:

- ▶ Rack Model 7014-T00
- ▶ Rack Model 7014-T42
- ▶ Rack-Mounted Flat Panel Console Model 7316-TF2
- ▶ Uninterruptible Power System (UPS)

7.1 Model 7014-T00 overview

The IBM 7014 RS/6000 Rack Model T00 is a 1.8 meter 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. Its features include:

- ▶ 36 EIA units (36 U) 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

Figure 7-1 shows a 7014-T00 rack populated with pSeries 610 Model 6C1 servers.

Figure 7-1 7014-T00 with pSeries 610 servers



7.2 Model highlights 7014-T00

The following describes the highlights of the 7014-T00 rack.

- ▶ Taller, stronger 19-in. wide rack provides increased height, weight, and power distribution capacity for RS/6000 systems.
- ▶ 1.8 meter Model T00 capacity is 36 U (EIA units).
- ▶ Optional features include front door, removable side panels, and side to side joining of multiple racks into suites.
- ▶ Optional ruggedized rack feature provides added earthquake protection with modular rear brace, concrete floor bolt down hardware, and bolt-in steel front filler panels.
- ▶ Available in standard black or optional white color.
- ▶ Model T00 supports both AC and DC configurations.

7.2.1 Description

The 1.8 meter 7014-T00 rack has 36 EIA units (36U) of vertical mounting space for 19-in. mounting-width drawers. An EIA unit is an industry standard for indicating vertical mounting space in a rack, where 1 U equals 44.45 mm (1.75 in.). This rack meets 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 are also orderable in a white color that matches the Model R00 rack and other older RS/6000 models.

7.2.2 Side panels and side to side rack suite attachment

Unlike previous RS/6000 racks, the rack side panels are optional, and can be quickly installed or removed from the outside the rack, with no tools required. This design allows a row of racks without side panels to be bolted together in a continuous suite, using the optional side to side rack connecting hardware. When multiple racks are joined in this way, cables can be easily 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. The trim pieces cover the space between each rack for an enhanced appearance and for additional protection of the equipment inside the racks. Side panels are needed only for the two end racks of the suite.

7.2.3 Front door and front trim kit

The optional front door has an appearance that is consistent with the latest pSeries Enterprise Systems. The main surface is steel, with a perforated hole pattern that is designed to provide the required ventilation and sufficient visibility of indicator lights for equipment installed inside the rack. Each rack must be ordered with either a front door or a front trim kit (but not both). The front trim kit consists of steel parts that snap into the front of the rack at the left, right, and top edges. They are painted to match the overall rack color and present an attractive appearance that is consistent with the RS/6000 brand image.

7.2.4 Ruggedized rack feature

For enhanced rigidity and stability of the rack, the optional Ruggedized Rack Feature provides additional hardware that reinforces the rack and anchors it to the floor. This hardware is designed primarily for use in locations where earthquakes are a concern. The feature includes a large steel brace that bolts into the rear of the rack. It is hinged on the left side so it can swing out of the way for easy access to the rack drawers when necessary. The Ruggedized Rack Feature also includes hardware for bolting the rack to a concrete floor, and bolt-in steel filler panels for any unoccupied spaces in the rack.

7.2.5 Other standard equipment and optional features

Each rack is shipped with standard equipment that includes a rear door, front and rear anti-tilt plates, four leveling feet, snap-in filler panels for any unoccupied rack space, and an AC power distribution unit (PDU) mounted inside the rack in an area that is separate from the drawer mounting area. There is room for up to 4 PDUs without using any of the rack's 36 U capacity. The rack drawer mounting area is centered relative to the two sides of the rack, which provides ample room to route cables within the frame on both 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 may be ordered instead of the standard AC PDU.

7.2.6 Physical specifications, operating environment, and power

For physical specifications, operating environment, and power requirements, go to Appendix A, "Site and hardware planning information" on page 747.

AC power distribution bus

This section describes the two types of AC power distribution buses:

Type 6 power distribution buses have six IEC320-C13, 200V to 240V AC outlets. There are an additional two front outlets with limited access. The input AC power to the bus is not switched, so each outlet has a separate circuit breaker to protect against excessive current. Type 6 power distribution bus feature codes are 9171, 9173, 9174, and 6xxx. Figure 7-2 show the rear view of a 7014-T00 rack with a type 6 power distribution.




Figure 7-2 7014-T00 rack with AC power distribution bus type 6

Type 7 power distribution buses have nine IEC320-C13, 200V to 240V AC outlets. There are also two IEC320-C19, 200V to 240V AC outlets. The input AC power to the bus is not switched, so each group of three IEC320-C13, 200V to 240V AC outlets has a separate circuit breaker to protect against excessive current. Each IEC320-C19, 200V to 240V AC outlet has a separate circuit breaker. Type 7 power distribution bus feature codes are 9176, 9177, 9178, and 7xxx. Figure 7-3 on page 314 show a type 7 power distribution bus.

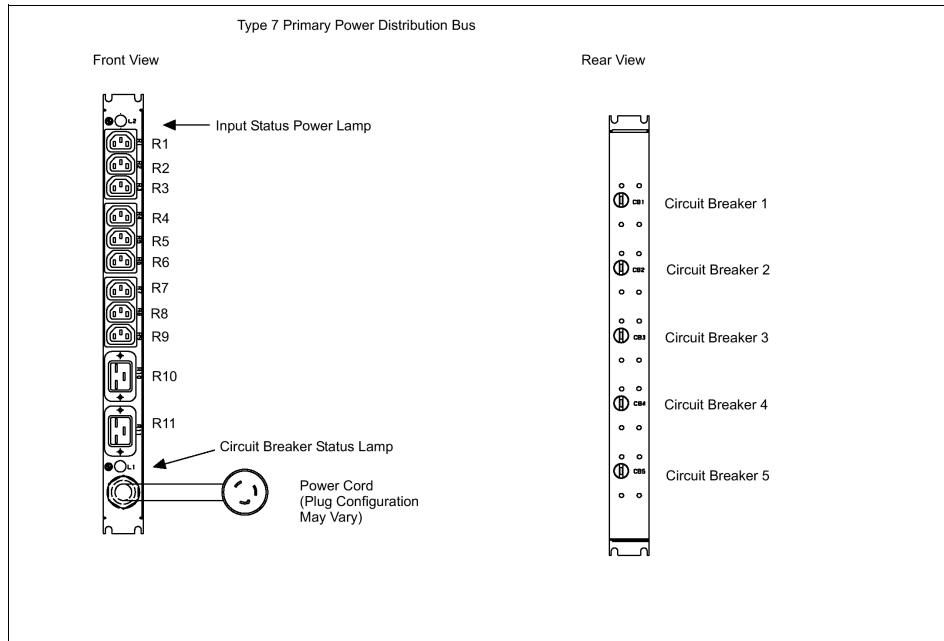


Figure 7-3 Type 7 power distribution bus

Note: The AC power requirements for the 7038 pSeries 650 server are significantly higher than previous IBM rack mounted servers and require the use of the new type 7 PDU (FC 9176, 9177, 9178, and 7xxx).

A rack can contain up to four vertically mounted AC power distribution buses. Two additional power distribution buses (each consuming 1 EIA location) may be mounted horizontally in the bottom rear of a 14T/0 rack and three additional power distribution buses (each consuming 1 EIA location) may be mounted horizontally in the bottom rear of a 14T/4 rack. The four vertical mounting slots will be filled first before consuming EIA locations in the rack. Type 6 AC power distribution buses contain eight 200V to 240V AC outlets to provide power to the devices and drawers. For type 6, the input AC power to the bus is not switched, so each outlet has a separate circuit breaker to protect against excessive current.

Limitations

The next section describe the limitations of the 7014-T00 rack.

- ▶ Each T00 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 (FC 6171 or FC 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 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 of the following features (FC 6171, FC 6173, FC 9171, or FC 9173) for each 9910-U33 installed.
- ▶ A maximum of 36 EIA units are available for mounting drawers in the rack.
- ▶ Maximum Weight of Drawers is 572 kg (1260 lb).
- ▶ A DC Power Distribution Panel (PDP), FC 6115, FC 6116, or FC 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.
- ▶ The T00 rack is designed to hold up to four Power Distribution Units vertically mounted in the back of the rack, two each on the left and right sides. Up to two additional PDUs can be installed horizontally in the rear of the rack, in the space that is normally used for rack-mounted drawers. Each horizontally

mounted PDU occupies 1 U of space in the rack, and therefore reduces the space available for mounting servers and other components.

- ▶ If PDU features 7177 or 9177 are ordered, the Power Cord Specify feature must be 9823.

Hardware requirements

The following provides an overview of supported PDU Combinations for 7014 Models T00 and T42.

Each 7014-T00 or 7014-T42 rack is shipped with one of the available base PDU features. The choice of which initial PDU is ordered for the rack will determine which additional PDU features may be ordered for the rack, either in the original rack order or in a later MES order. Table 7-1 shows the allowable PDU combinations within a single T00 or T42 rack.

Table 7-1 Allowable PDU combinations within a single T00 or T42 rack

Base PDU	Allowable Optional/Additional PDUs					
	6171	6173	6174	7176	7177	7178
9171	X			M ¹	M ²	
9173		X				M ³
9174			X			M
9176				X		
9177					X	
9178						X

Notes:

- ▶ X: Supported combination for initial or MES orders
- ▶ M: Supported combination for MES orders only
- ▶ ¹: For all line cords except FC 9823
- ▶ ²: Required for line cord FC 9823
- ▶ ³: 3-phase 16A (not 2 of 3 phase and 32A, such as FC 9173)

7.2.7 Publications

The following publication is shipped with the display product. Additional copies are available immediately:

- ▶ *7014 Series Model T00 and T42 Rack Installation and Service Guide*, SA38-0577

7.3 Model 7014-T42 overview

The IBM 7014 RS/6000 Rack Model T42 is a 2.0 meter rack that will address the special requirements of customers who want a very tall enclosure to house the maximum equipment in the smallest possible floor space. The T42 rack provides 42 EIA units (42 U) of usable space. Because of its height, special shipping and handling procedures may be required in some cases. 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:

- ▶ 42 EIA units (42 U) 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

Figure 7-4 shows a 7014-T42 rack.

Figure 7-4 7014-T42 rack



7.4 Model highlights 7014-T42

The following describe the highlights of the 7014-T42 rack.

- ▶ Taller, stronger 19-in. wide racks provide increased height, weight, and power distribution capacity for RS/6000 systems.
- ▶ 1.8 meter Model T00 capacity is 36 U (EIA units). 2.0 meter Model T42 capacity is 42 U.
- ▶ Optional features include front door, removable side panels, and side to side joining of multiple racks into suites.
- ▶ Optional ruggedized rack feature provides added earthquake protection with modular rear brace, concrete floor bolt down hardware, and bolt-in steel front filler panels.
- ▶ Available in standard black or optional white color.
- ▶ Model T42 supports AC only.

7.4.1 Description

The 2.0 meter 7014-T42 rack has 42 U of vertical mounting space for 19-in. drawers. An EIA unit is an industry standard for indicating vertical mounting space in a rack, where 1 U 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 are also orderable in a white color that matches the Model R00 rack and other older RS/6000 models.

7.4.2 Side panels and side to side rack suite attachment

Unlike previous RS/6000 racks, the rack side panels are optional, and can be quickly installed or removed from the outside the rack, with no tools required. This design allows a row of racks without side panels to be bolted together in a continuous suite, using the optional side to side rack connecting hardware. When multiple racks are joined in this way, cables can be easily 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. The trim pieces cover the space between each rack for an enhanced appearance and for additional protection of the equipment inside the racks. Side panels are needed only for the two end racks of the suite. The same removable side panels are used for both the T00 and the T42 racks. On T42 racks, there are additional side covers for the portion of the rack that extends above 1.8 meters.

7.4.3 Front door and front trim kit

The optional front door has an appearance that is consistent with the S80 and other current RS/6000 Enterprise Systems. The main surface is steel, with a perforated hole pattern that is designed to provide the required ventilation and sufficient visibility of indicator lights for equipment installed inside the rack. Each rack must be ordered with either a front door or a front trim kit (but not both). The front trim kit consists of steel parts that snap into the front of the rack at the left, right, and top edges. They are painted to match the overall rack color and present an attractive appearance that is consistent with the RS/6000 brand image. Because the T42 rack is taller than the T00 rack, the front doors, front trim kits, and side to side rack connecting kits require different sets of features for the two models.

7.4.4 Ruggedized rack feature

For enhanced rigidity and stability of the rack, the optional Ruggedized Rack Feature provides additional hardware that reinforces the rack and anchors it to the floor. This hardware is designed primarily for use in locations where earthquakes are a concern. The feature includes a large steel brace that bolts into the rear of the rack. It is hinged on the left side so it can swing out of the way for easy access to the rack drawers when necessary. The Ruggedized Rack Feature also includes hardware for bolting the rack to a concrete floor, and bolt-in steel filler panels for any unoccupied spaces in the rack.

7.4.5 Other standard equipment and optional features

Each rack is shipped with standard equipment that includes a rear door, front and rear anti-tilt plates, four leveling feet, snap-in filler panels for any unoccupied rack space, and an AC power distribution unit (PDU) mounted inside the rack in an area that is separate from the drawer mounting area. There is room for up to 4 PDUs without using any of the rack's 42 U capacity. The rack drawer mounting area is centered relative to the two sides of the rack, which provides ample room to route cables within the frame on both the left and right sides. There are numerous anchor points formed into the frame where cables can be tied down for strain relief.

7.4.6 Physical specifications

For physical specifications, operating environment, and power requirements, go to Appendix A, "Site and hardware planning information" on page 747.

AC power distribution bus

This section describes the two types of AC power distribution buses.

Type 6 power distribution buses have six IEC320-C13, 200V to 240V AC outlets. There are an additional two front outlets with limited access. The input AC power to the bus is not switched, so each outlet has a separate circuit breaker to protect against excessive current. Type 6 power distribution bus feature codes are 9171, 9173, 9174, and 6xxx. Figure 7-5 on page 321 show the rear view of a 7014-T42 rack with a type 6 power distribution.




Figure 7-5 Rear view of a 7014-T42 rack with a type 6 power distribution

Type 7 power distribution buses have nine IEC320-C13, 200V to 240V AC outlets. There are also two IEC320-C19, 200V to 240V AC outlets. The input AC power to the bus is not switched, so each group of three IEC320-C13, 200V to 240V AC outlets has a separate circuit breaker to protect against excessive current. Each IEC320-C19, 200V to 240V AC outlet has a separate circuit breaker. Type 7 power distribution bus feature codes are 9176, 9177, 9178, and 7xxx. Figure 7-6 on page 322 show a type 7 power distribution bus.




Figure 7-6 Type 7 power distribution bus

Note: The AC power requirements for the 7038 pSeries 650 server are significantly higher than previous IBM rack mounted servers and require use of the new type 7 PDU (FC 9176, 9177, 9178, and 7xxx).

A rack can contain up to four vertically mounted AC power distribution buses. Two additional power distribution buses (each consuming 1 EIA location) may be mounted horizontally in the bottom rear of a 14T/0 rack and three additional power distribution buses (each consuming 1 EIA location) may be mounted horizontally in the bottom rear of a 14T/4 rack. The four vertical mounting slots will be filled first before consuming EIA locations in the rack. Type 6 AC power distribution buses contain eight 200V to 240V AC outlets to provide power to the devices and drawers. For type 6, the input AC power to the bus is not switched, so each outlet has a separate circuit breaker to protect against excessive current.

Limitations

The next section describe the limitations of the 7014-T42 rack.

- ▶ 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 (FC 6171 or FC 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 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 of the following features (FC 6171, FC 6173, FC 9171, or FC 9173) for each 9910-U33 installed.
- ▶ A maximum of 42 EIA units are available for mounting drawers in the rack.
- ▶ Maximum Weight of Drawers is 667 kg (1470 lb).
- ▶ The T42 rack requires special handling because of its height. If a system is ordered with a 42 U rack, it cannot be shipped by air, and two additional days will be required for shipping. Requested arrival dates should be adjusted accordingly, that is, the system will typically arrive two days later than the calculated arrival date (CAD). Site planning should be done in advance to assure that the rack will fit under any low doorways at the customer location.

- ▶ 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 T42 rack is designed to hold up to four Power Distribution Units vertically mounted in the back of the rack, two each on the left and right sides. Up to three additional PDUs can be installed horizontally in the rear of the rack, in the space that is normally used for rack-mounted drawers. Each horizontally mounted PDU occupies 1 U of space in the rack, and therefore reduces the space available for mounting servers and other components.
- ▶ If PDU features 7177 or 9177 are ordered, the Power Cord Specify feature must be 9823.

Hardware requirements

The following provides an overview about supported PDU Combinations for 7014 Models T00 and T42.

Each 7014-T00 or 7014-T42 rack is shipped with one of the available base PDU features. The choice of which initial PDU is ordered for the rack will determine which additional PDU features may be ordered for the rack, either in the original rack order or in a later MES order. Table 7-2 shows the allowable PDU combinations within a single T00 or T42 rack.

Table 7-2 Allowable PDU combinations within a single T00 or T42 rack

	Allowable Optional/Additional PDUs					
Base PDU	6171	6173	6174	7176	7177	7178
9171	X			M ¹	M ²	
9173		X				M ³
9174			X			M
9176				X		
9177					X	
9178						X

Notes:

- ▶ X: Supported combination for initial or MES orders
- ▶ M: Supported combination for MES orders only
- ▶ ¹: For all line cords except FC 9823
- ▶ ²: Required for line cord FC 9823

- ▶ ³: 3-phase 16A (not 2 of 3 phase and 32A such as FC 9173)

7.5 Model 7316-TF2 overview

The IBM 7316 Model TF2 is a rack-mounted, flat panel console kit consisting of a Flat Panel Color Monitor, Rack Tray for monitor, keyboard, optional VGA switch, and mounting brackets.

Figure 7-7 show a 7316-TF2 flat panel console.

Figure 7-7 7316-TF2 flat panel console



7.5.1 Model highlights 7316-TF2

The IBM 7316-TF2 rack-mounted Flat Panel Console offers:

- ▶ Slim, sleek, lightweight design for flexibility in situations where space constraints are a critical issue
- ▶ A TFT monitor with bright, clear, and pure images on a flat screen, without the glare of most CRTs, and truly accurate images with virtually no distortion
- ▶ The ability to mount the monitor in a 19-in. standard rack, taking up only 1 U (1.75 in.) of rack space
- ▶ The IBM Space Saver 2 Keyboard, which mounts in the 7316-TF2 rack keyboard tray
- ▶ Support for the 2 x 8 VGA Switch

7.5.2 Description

The 7316-TF2 is a solution for attaching a monitor to a rack-mounted system. Its space-saving design occupies only 1 U (1.75-in.) in a rack and the monitor will mount in a 19-in. standard rack drawer, using the flat panel monitor rack-mount kit. The Model TF2 replaces the Model TF1, a 3 U flat panel console solution.

The flat panel color monitor offers advanced display characteristics and a stylish, space-saving advantage. This monitor combines a bright 15-in. viewable image and 1024 x 768 display resolution with a space-saving package design. Using thin film transistor (TFT) LCD technology, the monitor features a 304.1 mm x 228.1 mm viewable area.

The IBM Space Saver 2 Keyboard, specifically designed for rack-mounted servers such as the IBM @server pSeries 630 Model 6C4, is a mandatory feature for the 7316-TF2.

The 7316-TF2 also supports the VGA switch (FC 4200) that allows users single-point access and control of multiple servers from one console.

Additional information on the VGA switch can be found at:

<http://www.avocent.com/web/en.nsf>

Under "Select a product", select **280ES**.

7.5.3 Physical specifications

This section describes the physical specifications of the IBM 7316-TF2 rack-mounted Flat Panel Console:

- ▶ Weight: 37 lbs.
- ▶ Height: 1.75 inches (12" with monitor tray pulled out and TFT monitor opened)
- ▶ Width: 19 inches
- ▶ Depth: 29.6 inches

7.6 Uninterruptible Power System (UPS)

This section contains information about the following UPS systems:

- ▶ 9910-Axx (APC)
- ▶ 9910-Pxx (Powerware)
- ▶ 9910-Mxx (MGE)

IBM offers intelligent, online and line-interactive single phase Uninterruptible Power Systems (UPS) from American Power Conversion (APC), Powerware, and MGE.

The 9910 UPS Solutions are specifically designed for today's RS/6000/pSeries workstations, workstation servers, and enterprise servers. This cost-effective power protection bundle provides premium power management for the RS/6000/pSeries environments. These offerings provide a single point of purchase for your IBM equipment and your Power Protection needs and a single point of contact for service if necessary. The 9910 comes with all the cables and software necessary for connecting and communicating to the IBM equipment. They also provide a more robust warranty package than you typically will see with other UPS offerings in the market today. They include next day element replacement and extended lengths of warranty protection (APC three years and Powerware/Best Power and MGE five years). Warranty service is provided through IBM-SERV and managed by IBM. All UPS solutions have been tested by IBM Server Group Power Systems Engineering specifically for the defined IBM hardware models ensuring coordinated operations to help maximize system availability.

Currently, the UPS solutions for RS/6000/pSeries include products from APC (9910-Axx models), Powerware (9910-Pxx models) and MGE (9910-Mxx models). Table 7-3 provides an overview about the several 9910 models.

Table 7-3 Overview about the several 9910 models

Model	Manufacturer	Type	VA	Voltage	Watts
A08	APC	BK500 CS 500VA	500	110-127	300
A10	APC	SU700Net 700VA	700	110-127	450
A13	APC	Smart-UPS RS 1500VA	1400	110-127	980
A30	APC	Smart-UPS 5000VA rack mount	5000	200-240	3750
P07	Powerware	5115 750VA	750	110-127	500
P18	Powerware	5125 1500i rackmount	1500	154-288	1340
P22	Powerware	5125 2200VA	1920	110-127	1600
P33	Powerware	5125 rackmount	3000	208-240	2700
P60	Powerware	5140 EXT rackmount	6000	200-240	6000
M10	MGE	Pulsar Evolution 1100VA	1100	110-127	700
M15	MGE	Pulsar Evolution rackmount	1500	150-294	1000
M30	MGE	Pulsar Evolution rackmount	3000	150-294	2000

7.6.1 Rack rules

The appropriate UPS is determined by the sum of the watts assigned to the pSeries rack-mount drawers. In general, for the T00 and T42 racks, if the rack specify codes assigned to the pSeries drawers or I/O in a single rack totals two or less, the P18 or M15 is functional. If the number of rack specify codes is more than two, an additional P18 or M15 can be ordered for each multiple of two specify codes. For example, two rack specify codes order either one P18 or one M15, four rack specify codes order either two P18s, two M15s, or a larger rack UPS (P33, M30, A30, or P60). Do not include the 9910 specify code in the count. A maximum of either four P18s and four FC 6608 or eight M15s can be ordered per rack. A total UPS specify code count cannot exceed six. For example, two P18s and four 6608's would be the maximum UPS configuration per rack.

The P18 or M15 is a perfect solution for rack customers who will only have a two specify code rack configuration or want the flexibility to purchase UPS protection when they grow their rack configuration requirements.

For the T00 and T42 racks, if the number of rack specify codes are four or less, the P33, M30, A30, and P60 are functional. If the number of specify codes are more than four, an additional P33 or M30 can be ordered for each multiple of four specify codes. For example, four rack specify codes order one P33 or one M30, eight specify codes order two P33s or two M30s or a larger rack UPS (A30 or P60). Do not include the 9910 specify codes in the count. A maximum of either two P33s and four 6607s or two M30s and four 3401s can be ordered per rack. For all T00 and T42 racks with p650s or for all cases where only one 9910 UPS is desired, then the A30 or P60 should be used. For more information, including graphs/charts for the Powerware or MGE product configuration for IBM racks, refer to the following Web sites:

- ▶ For Powerware

<http://www.oem.powerware.com/ibm-ups/9910solutions.html>

- ▶ For MGE

<http://extranet.mgeups.com/accounts/ibm>

7.6.2 Customer setup (CSU)

All RS/6000 UPSs ordered as stand-alone units are customer setup. All deskside/desktop UPS models are also customer setup. 9910 rack-mounted UPSs are plant installed if ordered with the RS/6000 system. If rack-mounted UPS is ordered as a stand-alone (MES), the customer is responsible for installation.

7.6.3 Devices supported

The following link give information about the supported devices.

For Powerware, see:

<http://www.oem.powerware.com/ibm-ups/pdf/rsdream.pdf>

For MGE, see:

<http://extranet.mgeups.com/accounts/ibm>

7.6.4 Battery runtime

This section provide information about battery runtime of the different 9910 models.

For 9910-Axx (APC)

To see the battery runtime for 9910-Axx (APC) UPS systems, see the following links. To convert the IBM model name to the APC model name, use Table 7-3 on page 327.

- ▶ For Model 9910-A08, APC Type BK500:
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=17
- ▶ For Model 9910-A10, APC Type SU700NET:
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=165
- ▶ For Model 9910-A13, APC Type Smart-UPS RS 1500VA
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=23
- ▶ For model 9910-A30,APC Type SU5000RMT5U
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=165

For 9910-Pxx (Powerware)

To see the battery runtime for 9910-Pxx (Powerware) UPS systems, see the following link:

<http://www.oem.powerware.com/ibm-ups/pdf/rsdream.pdf>

For 9910-Pxx (MGE)

To see the battery run time for 9910-Mxx (MGE) UPS systems, see the following link:

<http://extranet.mgeups.com/accounts/ibm/doc/xlsizing.pdf>

7.6.5 UPS Model 9910-Axx

This section describes the following 9910-Axx models:

- ▶ 9910-A08 tower model
- ▶ 9910-A10 tower model
- ▶ 9910-A13 tower model
- ▶ 9910-A30 rack model

The 9910 Axx models for RS/6000/pSeries are products from APC. Additional information on these models and 9910 Axx models that have been withdrawn can be found at (for model 9910-Axx (APC)):

<http://sturgeon.apcc.com/sizing.nsf/search>

Search for the argument (IBM Smart UPS) or the APC type.

Choosing the right APC UPS

Table 7-4 provides application recommendations for the APC backups.

Table 7-4 APC application matrix

	9910-A08	9910-A10	9910-A13	9910-A30
7043-43P 150	X	X	X	X
7044-44P 170	X	X	X	X
7044-44P 270		X	X	X
7028-6E1 p610		X	X	X
7028-6C1 p610			X	X
7025-6F0 p620			X	X
7025-6F1 p620			X	X
7028-6E4 p630			X	X
7028-6C4 p630			X	X
7029-6E3 p615		X	X	X
7029-6C3 p615			X	X
7026-B80 p640			X	X
7046-B50			X	X
7026-6M1, 6H1, 6H0 p660			X	X

	9910-A08	9910-A10	9910-A13	9910-A30
9112 265		X	X	X
9114 275		X	X	X

9910-A08

The following section provides information for the Model 9910-A08.

Highlights

The following lists describes the highlights of the Model 9910-A08:

- ▶ User-replaceable, hot-swappable batteries.
- ▶ Double-Boost and SmartTrim correct over and under voltages without draining the battery.
- ▶ Informative LED front panel display.
- ▶ PowerChute plus for AIX for advanced UPS power management, graceful shutdown, battery replacement warnings, and diagnostics.

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-A10:

- ▶ Width: 9.14 cm (3.6 in.)
- ▶ Depth: 28.45 cm (11.2 in.)
- ▶ Height: 16.51 cm (6.5 in.)
- ▶ Weight: 6.27 kg (13.8 lb)
- ▶ Temperature: 0° C to +40° C (32° F to 104° F)
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: <45 dB of normal mode
- ▶ Electrical power: 110-120 volts 50 or 60 hertz

9910-A10

The following section provides information for the Model 9910-A10.

Highlights

The following list describes the highlights of the Model 9910-A10:

- ▶ User-replaceable, hot-swappable batteries.
- ▶ Double-Boost and SmartTrim correct over and under voltages without draining the battery.

- ▶ Informative LED front panel display.
- ▶ PowerChute plus for AIX for advanced UPS power management, graceful shutdown, battery replacement warnings, and diagnostics.

Physical specifications and operating environment

The following lists gives the physical specifications for the operating environment for the Model 9910-A10:

- ▶ Width: 13.72 cm (5.4 in.)
- ▶ Depth: 35.81 cm (14.1 in.)
- ▶ Height: 15.75 cm (6.2 in.)
- ▶ Weight: 14.55 kg (30 lb)
- ▶ Operating Temperature: 0° C to +40° C (32° F to 104° F)
- ▶ Storage Temperature: -15 to +45° C (+5 to +113° F)
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: <41 dB of normal mode
- ▶ Electrical power: 110-120 volts 50 or 60 hertz

9910-A13

The following section provides information for the Model 9910-A13.

Highlights

The following list describes the highlights of the Model 9910-A13.

- ▶ User-replaceable, hot-swappable batteries.
- ▶ Double-Boost and SmartTrim correct over and under voltages without draining the battery.
- ▶ Informative LED front panel display.
- ▶ PowerChute plus for AIX for advanced UPS power management, graceful shutdown, battery replacement warnings, and diagnostics.

Physical specifications and operating environment

The following list gives the physical specifications the operating environment for the Model 9910-A13.

- ▶ Width: 170 mm (6.7 in.)
- ▶ Depth: 439 mm (17.3 in.)
- ▶ Height: 216 mm (8.5 in.)
- ▶ Weight: 24.1 kg (53 lb)

- ▶ Temperature: 0° C to +40° C (32° F to 104° F)
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: <45 dB of normal mode
- ▶ Electrical power: 110-120 volts 50 or 60 hertz

9910-A30

The following section provides information for the Model 9910-A30.

Highlights

The following list describe the highlights of the Model 9910-A30:

- ▶ 5 U vertical rack mount UPS
- ▶ APT's AP9606 SNMP/Web Management Card
- ▶ PowerChute plus for AIX and cables
- ▶ PowerChute Network Shutdown for AIX
- ▶ Additional APC SmartSlot Accessory Bays

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-A30.

- ▶ Width: 483.0 mm (19 in.)
- ▶ Depth: 635.0 mm (25 in.)
- ▶ Height: 222.3 mm (38.75 in.)
- ▶ Weight: 102.2 kg (225 lbs)
- ▶ Operating Temperature: 0° C to +40° C
- ▶ Storage Temperature: -15 to +40° C
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Audible Noise at One Meter: <53 dBA
- ▶ Operation Elevation: 0 to +3000 m (0 to +10,000 ft.)
- ▶ Input Voltage Range: 157-255VAC

7.6.6 UPS Model 9910-Pxx

This section describes the following 9910-Pxx models:

- ▶ 9910-P07 tower model
- ▶ 9910-P18 rack model

- ▶ 9910-P22 tower model
- ▶ 9910-P33 rack model
- ▶ 9910-P60 rack model

The 9910 Pxx models for RS/6000/pSeries are products from Powerware. Additional information on these models and 9910 Pxx models that have been withdrawn can be found at (for model 9910-Pxx (Powerware)):

<http://www.oem.powerware.com/ibm-ups/9910solutions.html>

Choosing the right Powerware UPS

Table 7-5 can help to choose the right Powerware UPS.

Table 7-5 Help to choose the right Powerware UPS

	9910-P07	9910-P18	9910-P22	9910-P33	9910-P60
7043-43P 150	X		X		
7044-44P 170, 270	X		X		
7028-6E1 p610	X		X		
7025-6F0 p620			X		
7025-6F1 p620			X		
7028-6E4 p630			X		
7046-B50		X		X	X
7028-6C1 p610		X		X	X
7028-6C4 p630		X		X	X
7026-B80 p640		X		X	X
7026-6M1, 6H1, 6H0 p660		X		X	X

	9910-P07	9910-P18	9910-P22	9910-P33	9910-P60
9112 265, 9114 275	X		X		

9910-P07

The following section provides information for the Model 9910-P07.

Highlights

The following list describes the highlights of the Model 9910-P07:

- ▶ Advanced Battery Management (ABM) doubles battery service life.
- ▶ Buck and Boost voltage regulation corrects incoming fluctuations.
- ▶ Network transient protector.
- ▶ Hot-swappable batteries simplify service.
- ▶ Bundled with Software Suite power management software.

Physical specifications and operating environment

The following list gives the physical specifications the operating environment for the Model 9910-P07:

- ▶ Width: 15.0 cm (5.9 in.)
- ▶ Depth: 33.5 cm (13.2 in.)
- ▶ Height: 19.4 cm (7.6 in.)
- ▶ Weight: 12.3 kg (27.3 lbs)
- ▶ Temperature: 10° C to +40° C
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: 45dBA at 1 meter
- ▶ Electrical power: nominal settings 110, 120 VAC +/-20% nominal

9910-P18

The following section provides information for the Model 9910-P18.

Highlights

The following list describes the highlights of the Model 9910-P18:

- ▶ Advanced Battery Management (ABM) doubles battery service life.
- ▶ 2 U rack height conserves valuable rack space.
- ▶ Extended Battery Modules and two load segments prolong backup times.

- ▶ Buck and double boost voltage regulation with pure sine wave output.
- ▶ Hot-swappable batteries simplify service.
- ▶ Bundled with Software Suite power management software.

Physical specifications and operating environment

The following list gives the physical specifications the operating environment for the Model 9910-P18:

- ▶ Width: 48.0 cm (19 in.)
- ▶ Depth: 49.0 cm (19.4 in.)
- ▶ Height: 9.0 cm (3.5 in.)
- ▶ Weight: 15.4 kg (34 lbs)
- ▶ Temperature: 0° C to +40° C (32° F to 104° F)
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: <40 dBA typical
- ▶ Input Voltage Range: 154-288 V nominal Settings 220, 230, or 240 Vac. Configurable using a free-of-charge PW5125 Function Configurator, found at: <http://www.powerware.com>
 - a. Select **PRODUCTS**.
 - b. Select **POWERWARE 5125**.
 - c. Then select **PW 5125 FUNCTION CONFIGURATOR**.

9910-P22

The following section provides information for the Model 9910-P22.

Highlights

The following list describes the highlights of the Model 9910-P22:

- ▶ Advanced Battery Management (ABM) doubles battery service life.
- ▶ Two load segments allow scheduled shutdowns and prolong backup times.
- ▶ Buck and double boost voltage regulation with pure sine wave output.
- ▶ Hot-swappable batteries simplify service.
- ▶ Bundled with Software Suite power management software.

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-P22:

- ▶ Width: 20.2 cm (8 in.)

- ▶ Depth: 50.0 cm (19.4 in.)
- ▶ Height: 25.0 cm (9.8 in.)
- ▶ Weight: 32 kg (68.3 lbs)
- ▶ Temperature: 10° C to +40° C
- ▶ Relative Humidity: 0% to 95% noncondensing
- ▶ Noise Rating: 40dBA at 1 meter
- ▶ Electrical power: 77-153V nominal settings 110, 120, or 127 VAC

9910-P33

The following section provides information for the Model 9910-P33.

Highlights

The following list describes the highlights of the Model 9910-P33.

- ▶ 2U Rack-mounted UPS.
- ▶ Advanced line-interactive single phase UPS.
- ▶ ABM doubles battery service life and optimizes recharge time.
- ▶ 60 day notification of battery end of life.
- ▶ Load Segmentation for three separate receptacle groups.
- ▶ Multiport Serial Expander Card provides up to three server connections.
- ▶ Powerware Software Suite, which includes LanSafe III.

Physical specifications and operating environment

The following list gives the physical specifications the operating environment for the Model 9910-P33:

- ▶ Width: 44 cm (17.5 in.)
- ▶ Depth: 61 cm (24.0 in.)
- ▶ Height: 9 cm (3.5 in.)
- ▶ Weight: 82 lbs
- ▶ Temperature: 10° C to +40° C
- ▶ Relative Humidity: 5% to 90% noncondensing
- ▶ Noise Rating: <45dBA

9910-P60

The following section provides information for the Model 9910-P60.

Highlights

The following list describes the highlights of the Model 9910-P60:

- ▶ 6U rack-mount UPS.
- ▶ Manual bypass switch.
- ▶ Extended Battery Modules and 5 Load Segments prolong backup times.
- ▶ Hot-swappable battery modules.
- ▶ Power Factor Corrected design protects more equipment.
- ▶ Advanced Battery Management (ABM) doubles battery service life.
- ▶ Bundled with jumper cords/Software Suite for power management.

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-P60:

- ▶ Width: 17.25 in.
- ▶ Depth: 24.3 in.
- ▶ Height: 10.5 in.
- ▶ Weight: 250 lbs
- ▶ Temperature: 10° C to +40° C (50° F to 104° F)
- ▶ Relative Humidity: 5% to 95% noncondensing
- ▶ Noise Rating: <55 dBA typical

7.6.7 UPS Model 9910-Mxx

This section describes the following 9910-Pxx models:

- ▶ 9910-M10/M11 tower model
- ▶ 9910-M15 rack model
- ▶ 9910-M30 rack model

The 9910 Mxx models for RS/6000/pSeries are products from Powerware. Additional information on these models and 9910 Mxx models that have been withdrawn can be found at (for model 9910-Mxx (MGE)):

<http://extranet.mgeups.com/accounts/ibm/>

Choosing the right MGE UPS

Table 7-6 can help to choose the right MGE UPS.

Table 7-6 Choose the right MGE UPS

	9910-M10/M11	9910-M15	9910-M30
7043 -43P 150	X		
7044-44P 170, 270	X		
7028-6E1 p610	X		
7025-6F0 p620	X		
7025-6F1 p620	X		
7028-6E4 p630	X	X	
7046-B50			
7028-6C1 p610		X	
7028-6C4 p630		X	
7026-B80 p640		X	
7026-6M1, 6H1,6H0 p660		X	
7026-6H1 + 7133 D40 Storage Expansion			X

9910-M10/M11

The following section provides information for the Model 9910-M10/M11.

Note: The 9910-M10 and 9910-M11 are similar; only the voltage is different. The 9910-M10 is for North America (110V-127V); the 9910-M11 is for EMEA/Asia (200V-240V).

Highlights

The following list describes the highlights of the Model 9910-M10/M11:

- ▶ Receptacle control: Reboot up to two locked up systems remotely.
- ▶ Standard USB port.
- ▶ Cable locking system: IT devices cannot be unplugged accidentally.

- ▶ High Frequency technology: More compact and lightweight.
- ▶ Bundled with Power Management Software AIX, Linux, Windows, and more.

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-M10/M11:

- ▶ Width: 150 mm (5.91 in.)
- ▶ Depth: 415 mm (16.34 in.)
- ▶ Height: 237 mm (9.33 in.)
- ▶ Weight: 11.5 kg (25.3 lb)
- ▶ Temperature: 0° C to +35° C (32° F to 95° F)
- ▶ Relative Humidity: 0% to 90% Non-condensing
- ▶ Noise Rating: 40 dBA on Utility, 45 dBA on battery/buck and boost
- ▶ Electrical power 9910-M10: 70/80V -153V adjustable to 110, 120, or 127 VAC
- ▶ Electrical power 9910-M11: 150 V-294 V adjustable to 200, 220, or 240 Vac

9910-M15

The following section provides information for the Model 9910-M15.

Highlights

The following list describes the highlights of the Model 9910-M15:

- ▶ Receptacle control: Reboot up to two locked up systems remotely.
- ▶ Standard USB port.
- ▶ Cable locking system: IT devices cannot be unplugged accidentally.
- ▶ High Frequency technology: More compact and lightweight.
- ▶ Bundled with Power Management Software AIX, Linux, Windows, and more,

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-M15:

- ▶ Width: 438 mm (17.24 in.)
- ▶ Depth: 522 mm (20.55 in.)
- ▶ Height: 44.45 mm (1.75 in.) 1 U
- ▶ Weight: 19 kg (41.9 lb)
- ▶ Temperature: 0° C to +40° C (32° F to 95° F)

- ▶ Relative Humidity: 0% to 90% Non-condensing
- ▶ Noise Rating: 40 dBA on Utility, 45 dBA on battery/buck and boost
- ▶ Electrical power: 160-295 V adjustable to 150-294 using supplied UPS Driver software

9910-M30

The following section provides information for the Model 9910-M30.

Highlights

The following list describes the highlights of the Model 9910-M30:

- ▶ Rack Mount 1 U in size.
- ▶ Receptacle control: Reboot up to two locked up systems remotely.
- ▶ Standard USB port.
- ▶ Cable locking system: IT devices cannot be unplugged accidentally.
- ▶ Rear bracket system: Installed in rack prior to shipment.
- ▶ High Frequency technology: More compact and lightweight.
- ▶ Bundled with Power Management Software AIX, Linux, Windows, and more.

Physical specifications and operating environment

The following list gives the physical specifications for the operating environment for the Model 9910-M30.

- ▶ Width: 438 mm (17.24 in.)
- ▶ Depth: 640 mm (25.2 in.)
- ▶ Height: 89.9 mm (3.5 in.) 2 U
- ▶ Weight: 36.6 kg (80.5 lb)
- ▶ Temperature: 0° C to +40° C (32° F to 95° F)
- ▶ Relative Humidity: 0% to 90% Non-condensing
- ▶ Noise Rating: 40 dBA on Utility, 45 dBA on battery/buck and boost
- ▶ Electrical power: 160-295 V adjustable to 150-294 using supplied UPS Driver software



High density cluster servers

This chapter discusses the IBM 7039-651 @server pSeries 655 and the Large Scale Servers - RS/6000 SP Systems.

8.1 7039-651 IBM @server pSeries 655 Server

The IBM @server pSeries 655 server continues IBM POWER4+ chip technology in a building-block approach to the world of high-performance, clustered technical and commercial computing. With the speed advantages provided by the powerful POWER4+ processor and its associated system architecture, fast system bus, memory, and input/output (I/O) subsystems, the pSeries 655 (p655) today provides the greatest sustainable memory bandwidth per fully populated rack than any other non-vector, clustered servers.

8.1.1 The pSeries 655 as a Cluster Server

The pSeries 655 delivers power, density, and flexibility. The heart of the rack-mounted pSeries 655 server is the four-chip POWER4+ multichip module (MCM), the basic building block for the pSeries 690 8-way to 32-way servers. The four-chip MCMs come in the two versions, the two-processor core per chip (8-way) feature operating at 1.5 GHz and the single-processor core per chip (4-way) feature operating at 1.7 GHz.

Both the 4-way and 8-way pSeries 655 servers are optimized for technical workloads, such as those found in engineering and scientific environments, as well as commercial business intelligence (data warehousing, data mining, extract, transformation and load, backup/restore operations, and so on). The servers using the 4-way MCM are optimized for memory bandwidth-intensive applications, such as computational fluid dynamics and structural analysis. A fabric bus interconnects the pSeries 655 servers MCM processor chips and operates at one-half the processor frequency. It supports a peak bandwidth of 51.2 GBps for the 8-way 1.5 GHz feature and 51.2 GBps for the 4-way 1.7 GHz feature (Figure 8-1 on page 345).

The pSeries 655 memory subsystem consists of Level 1 (L1), Level 2 (L2), Level 3 (L3) caches and main memory. L1 instruction cache and data cache are included in each processor core. Each processor chip has a 1.5 MB L2 cache onboard that operates at chip frequency.

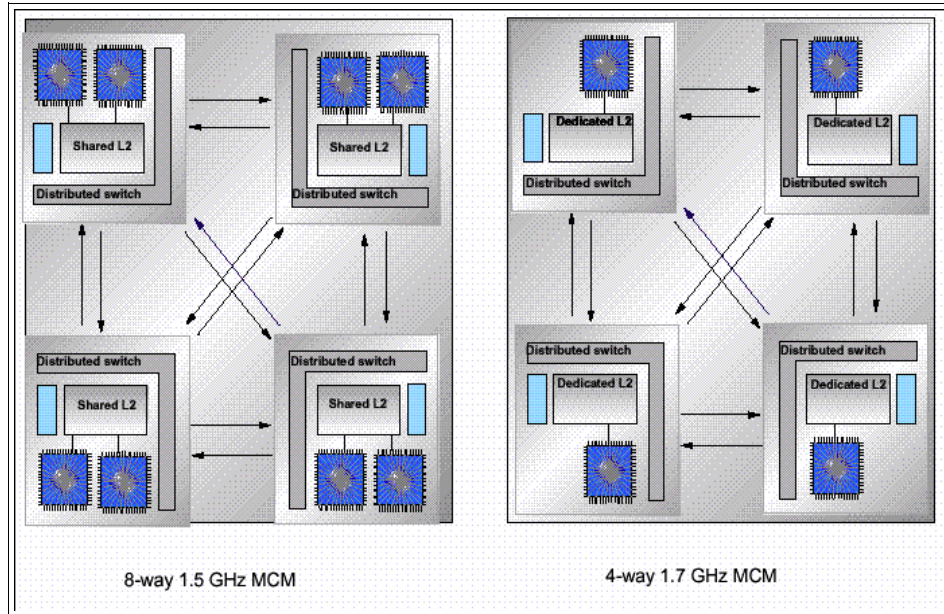


Figure 8-1 8-way 1.5 GHz and 4-way 1.7 GHz MCMs

On the 8-way MCM, the two cores on each processor chip share that chip's L2 cache, while on the 4-way MCM, each core has a dedicated L2 cache. A 32 MB L3 cache is located between each processor chip and main memory and operates at one-third of the chip frequency. The 4-way MCM has the full complement of L2 and L3 caches and therefore provides twice the L2 and L3 cache per processor as the 8-way MCM does. For pSeries 655 systems configured with two or four memory cards, the memory and L3 caches are interleaved at cache line granularity across the four memory controllers in the system. This creates a single "flat" shared memory with extraordinary peak and sustained bandwidth capability. The pSeries 655 contains four memory slots and the customer can select from between 4 GB, 8 GB, or 16 GB memory cards. The amount of main memory available, from 4 GB to 64 GB, will depend on the number and size of memory cards selected to populate the pSeries 655's four memory slots. The following memory configurations are offered (Table 8-1).

Table 8-1 Recommended p655 memory configurations

Total Memory	Configuration
4 GB	4 GB card in slot 1
8 GB	4 GB cards in slots 1 and 3
16 GB	4 GB cards in slots 1, 2, 3, and 4

Total Memory	Configuration
16 GB	8 GB cards in slots 1 and 3
32 GB	16 GB cards in slots 1 and 3
64 GB	16 GB cards in slots 1, 2, 3, and 4

Important: Although the pSeries 655 can be configured with one memory card, its best performance will be obtained when all four memory slots are filled with memory cards of the same size.

Peak bi-directional bandwidth between the L3 caches and the MCM and between the L3 caches and the memory controller is 64 GBps for the 8-way 1.5 GHz system. The corresponding value is 72.5 GBps for the 4-way 1.7 GHz system. Peak bi-directional bandwidth between the memory controller and main memory is 51.2 GBps on either system.

The pSeries 655 provides input/output (I/O) through the GX bus to two remote I/O (RIO-2) ports. The GX bus supports 4.0 GBps full duplex sustained bandwidth. The first RIO-2 port supports the service processor, two Ethernet ports, an integrated SCSI adapter, and three hot-plug/blind-swap PCI-X slots on the system board. The integrated SCSI adapter supports two hot-swappable, front-loadable disk drives on the p655. At present, the customer can select from 18.2 GB, 36.4 GB, 73.4 GB, and 146.8 GB drives. The second RIO-2 port provides a connection for an optional IBM 7040-61D I/O Drawer, which provides 20 hot-plug/blind-swap PCI slots and can contain up to 16 hot-swappable disk drives. RIO-2 drawers support a sustained bandwidth of 1.0 GBps in each direction.

8.1.2 The p655 Clustered System

IBM @server Cluster 1600 is a highly scalable cluster solution that includes pSeries servers and a choice of cluster management software, the Parallel Systems Support Programs for AIX Version 3.4 or 2.5 or Cluster System Management for AIX Version 1.3. Up to 64 individual pSeries 655 servers can be clustered together to support large-scale computational modeling and multi-terabyte parallel databases. A PSSP-managed cluster can optionally include an SP Switch2 high-performance cluster interconnect. PSSP-managed clusters can also include optional software for parallel application development, workload balancing, or distributed file system management. The p655 can be mixed and matched in a Cluster 1600, which includes other pSeries servers or RS/6000 SP nodes. Software components for the pSeries 655 server cluster include the AIX 5L Version 5.1 and 5.2 operating system and the Parallel System Support Programs (PSSP) for AIX Versions 3.4 or 3.5. PSSP provides a single

point-of-control for the cluster. Optional cluster software, including Parallel Environment, the Engineering and Scientific Subroutine Library (ESSL), Parallel ESSL, LoadLeveler®, and the General Parallel File System (GPFS), is also supported in a PSSP-managed cluster.

Figure 8-2 shows a cluster of servers managed by a control workstation.




Figure 8-2 Clustered system with p655 servers, Control Workstation, and HMC

Since the p655 is a 4 U high half-drawer wide system, sixteen pSeries 655 servers can be installed in an IBM 7042-W42 system frame, which is 30.9 inches wide by 56.8 inches deep and has a footprint of 12.2 square feet. A Cluster 1600 system consisting of one system frame fully populated with 8-way 1.5 GHz pSeries 655 servers will have 128 processors; a system frame fully populated with 4-way 1.7 GHz pSeries 655 servers will have 64 processors. Up to 64 pSeries servers, all 64 of which can be pSeries 655 servers, and 128 logical partitions (LPARs), can be interconnected in a Cluster 1600.

The Hardware Management Console (required for both PSSP and CSM-managed clusters), shown in Figure 8-2, provides asynchronous attachment to the Central Electronics Complex (CEC) to manage system control software, manage LPARs and control clusters, as well as to serve as the service

focal point. The Control Workstation (required only for PSSP-managed clusters), also shown in Figure 8-2, is used for systems management functions, such as installation and maintenance of system and application software, configuring volume groups, database setup and control, managing users, and other functions requiring root privileges.

A CSM-managed cluster cannot contain an SP Switch or SP Switch2. In addition, none of the optional cluster software listed above is available for a CSM-managed cluster.

8.1.3 Software requirements

AIX 5L Version 5.1 with Maintenance Level 3 is needed for the installation of the base operating system and the support of PSSP Version 3.5. Furthermore, APAR IY34495 for PSSP Version 3.4 or APAR IY34496 for PSSP Version 3.5 is required for support with PSSP. When running on a 64-bit kernel, only PSSP Version 3.5 is supported. Example 8-4 shows how PSSP recognizes the p655.

Figure 8-3 on page 349 shows the management software report on the cluster configuration.

```
[c179s][/]> splstdata -n 11 1 1
List Node Configuration Information

node# frame# slot# slots initial_hostname reliable_hostname dce_hostname default_route
processor_type processors_installed description on_switch primary_enabled LPAR_name
-----
-----
161 11 1 1 c59ih04.ppd.pok.i c59ih04.ppd.pok.i "" 9.114.213.125
MP 8 7039-651 1 true c59ih04

[c179s][/]> splstdata -b 11 1 1

List Node Boot/Install Information

node# hostname hdw_enet_addr srvr response install_disk last_install_image
last_install_time next_install_image lppsource_name pssp_ver selected_vg
-----
-----
161 c59ih04.ppd.pok.i 00096BE80041 0 disk hdisk0 bos.obj.node.aix51d
Fri_Sep_6_14:01:04 bos.obj.node.aix51d aix51d0212d PSSP-3.5 rootvg

[c179s][/]> splstdata -g 11 1 1
List Aggregate IP Database Information

node# adapt netaddr netmask hostname devicename update_interval
update_threshold
-----
-----
161 m10 9.114.213.28 255.255.255.192 c179san28.ppd.pok css0,css1 3 10
```

Figure 8-3 p655 in a PSSP Version 3.5 managed cluster

8.1.4 p655 Cluster considerations

The p655 can be integrated into a Cluster 1600 with the following configurations:

- ▶ A switchless cluster with the entire p655 as a single node.
- ▶ A switchless cluster with the p655 with two LPARs, each configured as a node.
- ▶ A cluster with a single or a double plane SP Switch2, where the p655 is off the switch and the p655 is a single node.
- ▶ A cluster with a single or a double plane SP Switch2, where the p655 is off the switch with both LPARs.
- ▶ A cluster with a single plane SP Switch2, where the p655 is on the switch using one SP Switch2 PCI-X Attachment Adapter (FC 8398) in one of its PCI-X slots and has no LPARs. Use slot 1.

- ▶ A cluster with a single plane SP Switch2, where the p655 is on the switch using one SP Switch2 PCI-X Attachment Adapter (FC 8398) in one of its PCI-X slots in one LPAR, while the other LPAR is off the switch. Use slot 1.
- ▶ A cluster with a single plane SP Switch2, where the p655 is on the switch using two SP Switch2 PCI-X Attachment Adapters (FC 8398), one in each LPAR.
- ▶ A cluster with a dual plane SP Switch2, where the p655 is on the switch using two SP Switch2 PCI-X Attachment Adapters (FC 8398). Use slots 1 and 3.
- ▶ A cluster with a dual plane SP Switch2, where the p655 is on the switch using two SP Switch2 PCI-X Attachment Adapters (FC 8398) in both PCI-X slots for one LPAR; the other LPAR is off the switch. Use slots 1 and 3.

8.1.5 7040-61D I/O Drawer

The pSeries 655 can be configured with a 7040-61D I/O Drawer. The 7040-61D I/O Drawer provides PCI adapter slots and internal disk capabilities for use with pSeries servers, such as the pSeries 655 and the p690/p670.

- ▶ Contains up to 16 hot-swappable disks structured into four 4-packs of 36.4 GB, 73.4 GB, or 146.8 GB disk drives; 2.3 TB maximum per I/O drawer. Separately powered backplanes with two integrated Ultra3 SCSI adapters per two 4-packs.
- ▶ 4U I/O drawer (split between two, or dedicated to one p655), up to five I/O drawers per frame (11.7 TB/frame). Each I/O drawer must be placed in the same frame as the server to which it is attached.
- ▶ PCI-X with RIO-2 connectivity, on PCI-X based planars available on the POWER4+ versions of the p655. Additional SP Switch2 connectivity with the SP Switch2 PCI-X Attachment Adapter (FC 8389) in the I/O Drawer.
- ▶ 20 PCI or PCI-X slots per drawer (p655 also has three internal slots, providing a total of 23 slots max per p655). Each 10 slots are powered separately.
- ▶ All ten PCI or PCI-X slots on each I/O drawer planar are capable of supporting either 64-bit or 32-bit PCI or PCI-X adapters. Each FC 6563 planar provides seven 3.3V, 66 MHz, 64-bit PCI slots and three 5V, 32 MHz, 64-bit PCI slots. Each FC 6571 planar provides ten PCI-X slots capable of supporting 3.3V signaling PCI or PCI-X adapters operating at speeds up to 133 MHz.
- ▶ Hot-plug slots permit PCI and PCI-X adapters to be added or replaced without extending the I/O drawer (blind-swap) while system remains available.

I/O Drawer attachment

The following are the configuration notes regarding the I/O drawer attachment:

- ▶ 7040-61D I/O drawers are always connected to the p655 CEC using RIO (Remote I/O) or RIO-2 loops. 61D connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. p655 systems with non-looped configurations could experience degraded performance and serviceability.
- ▶ RIO loop connections operate at 500 MHz. RIO loops connect to the 7039-651 CEC using RIO ports within the features FC 5511, FC 5513, FC 5515, or FC 5518.
- ▶ RIO-2 loop connections operate at 1 GHz. RIO-2 loops connect to the 7039-651 CEC using RIO-2 ports within the features FC 5515 and 5518.
- ▶ 7040-61D I/O drawers may be connected to the 7039-651 CEC in either single loop or dual loop mode. Dual loop mode is recommended whenever possible, as it provides the maximum bandwidth between the I/O drawer and the CEC.

8.1.6 Disk and boot devices

The following notes are concerned with the configuration of the disk and boot devices:

- ▶ A minimum of two internal SCSI hard disks are required per p655 server. It is recommended that these disks be utilized as mirrored boot devices. This configuration provides service personnel the maximum amount of diagnostic information if the system encounters errors in the boot sequence.
- ▶ Disk sizes offered are 18.2 GB (FC 3157), 36.4 GB (FC 3158), 73.4 GB (FC 3159), and 146.8 GB (FC 3275).
- ▶ Boot support is also available from local SCSI, SSA, and Fibre Channel adapters, or from the network using ENET or token-ring adapters.
- ▶ The p655 incorporates an Early Power Off Warning (EPOW) capability that assists in performing an orderly system shutdown in the event of a sudden power loss. IBM recommends use of the integrated battery backup features or an uninterruptible power system (UPS) to help ensure against loss of data due to power failures.
- ▶ No hot-swap disk capability exists within the p655 CEC, rather only within the I/O drawer.

8.1.7 7040-W42 System Rack

The 7040-W42 System Rack provides space for mounting system drawer components for pSeries servers, such as the pSeries 655 Central Electronics Complex (7039-651) and the 7040-61D I/O drawers. The Model W42 is a 24-in. rack that provides 42 U of rack space.

The 7040-W42 incorporates a 350V DC bulk power subsystem to provide power for the components within the rack. The bulk power subsystem incorporates redundant bulk power assemblies mounted in the front and rear sections of the top 8 U of the rack.

- ▶ Two Power Controller features (FC 6187) are required for all pSeries 655 systems, one per bulk power assembly. Each Power Controller provides power connections to support the systems four cooling fans. It also provides three additional connectors, which can be utilized to attach DC power converters contained in the CEC and each I/O drawer. Ten additional connector locations are provided by the Power Distribution Assembly (FC 6188). Two Power Distribution Assemblies (one per Bulk Power Assembly) may be added as required to provide more connections for large system configurations.
- ▶ Optional Integrated Battery Backup is available, if desired. The primary and secondary battery backup features (FC 6200 and FC 6201) interface with the power controllers of the front and rear bulk power assemblies. The battery backup features protect against power line disturbances and provide sufficient power to allow an orderly system shutdown in the event that the power sources fail. The battery backup features require up to 6 U of rack space and are always installed at the 13 U location of the primary 7040-W42 rack.
- ▶ p655 CECs and I/O drawers occupy specific locations in the 7040-W42 rack. Drawer placement is specified by the mounting location features ordered on the CEC or I/O drawer. The first p655 CEC is in Drawer #1. Additional drawers can have 7039-651 CECs, 7040-61D I/O, or IBF features of the 7040-W42 rack.

Rules for racking p655 components

CECs will be placed first at the left rack location (as viewed from the front), then the right rack location at the same EIA position. Both EIA positions will be filled before a new CEC EIA position can be filled.

At validation:

1. First, e-config will assign rack position features to all p655 systems that include a 61D with FC 3145.
2. e-config assigns a p655 that includes a 61D with three RIO-2 cables.

3. Second, e-config will:
 - a. Assign rack position features to a p655 system that includes a 61D with four RIO cables and without FC 3145.
 - b. Assign rack position features to a p655 system that does not include a 61D, but is connected to the second 61D planar. Iterate until all systems in this category are placed.
 - c. Assign rack position features to a p655 system that includes a 61D with two RIO cables and two RIO-2 cables.
 - d. Assign rack position features to a p655 system that includes a 61D with four RIO-2 cables.
4. Third, e-config will assign rack position features to p655 systems that include a 61D with two RIO cables and without feature 3145, or a 61D with two RIO-2 cables.
5. Fourth, e-config will assign rack position features to p655 systems that do not include a 61D, and are not connected to the second 61D planar using RIO or RIO-2 cables.

8.1.8 Feature codes on the p655

Table 8-2 lists the internal features that can be added to a p655 configuration. The status of a feature is indicative of these qualifications:

- X** Available. Indicates features that are available and are orderable on the specified model.
- W** Withdrawn. Indicates a feature that is no longer available to order.

Features not listed in the provided categories indicate that the feature is not supported on this model. Some categories, such as keyboards, mice, language specify codes, and power cords, are not included.

Table 8-2 Features used in the p655

Feature code	Description	Status
Host attachment adapter		
8398	SP Switch2 PCI-X Attachment Adapter	X
SCSI adapters		
6204	PCI Universal Differential Ultra SCSI Adapter	X
6203	PCI Dual Channel Ultra3 SCSI Adapter	X

Feature code	Description	Status
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter	X
5712	PCI-X Dual Channel Ultra320 SCSI Adapter	X
Serial adapters		
6231	128 MB DRAM Option Card	X
6230	Advanced SerialRAID Plus Adapter	X
6235	32 MB Fast-Write Cache Option Card	X
Fibre Channel adapters		
6228	2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus	W
6239	2 Gigabit Fibre Channel PCI-X Adapter	X
Graphics accelerators		
2848	POWER GXT135P Graphics Accelerator	X
2849	POWER GXT135P Graphics Accelerator with Digital Support	X
LAN adapters		
4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter	X
4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter	X
4961	IBM Universal 4-Port 10/100 Ethernet Adapter	X
4959	IBM Token-Ring PCI Adapter	X
2946	Turboways 622 Mbps PCI MMF ATM Adapter	X
5707	IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter	X
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	X
4962	10/100 Mbps Ethernet PCI Adapter II	X
5700	IBM Gigabit Ethernet-SX PCI-X Adapter	X
5701	IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter	X
WAN adapters		
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	X

Feature code	Description	Status
Adapter mounting hardware		
4599	PCI Blind Swap Cassette Kit, Single Wide Adapters, Universal	X
Miscellaneous		
4963	PCI Cryptographic Coprocessor (FIPS-4)	X
4960	IBM e-business Cryptographic Accelerator	X
2737	Keyboard/Mouse Attachment Card - PCI	X
Miscellaneous internal system		
6430	Service Support Processor	X
Asynchronous cables		
8121	Attachment Cable, Hardware Management Console to Host, 15-Meter	X
8120	Attachment Cable, Hardware Management Console to Host, 6-Meter	X
Fibre cables		
2456	LC-SC 50 Micron Fiber Converter Cable	X
2459	LC-SC 62.5 Micron Fiber Converter Cable	X
Graphics cables		
4242	6 Foot Extender Cable for Displays (15 pin D-shell to 15 pin D-shell)	X
Keyboard cables		
4256	Extender Cable - USB Keyboards, 2M	X
WAN cables		
2861	ARTIC960Hx 4-Port EIA-232 Cable	X
2864	ARTIC960Hx 4-Port V.35 (DTE) Cable	X
2863	ARTIC960Hx 4-Port X.21 Cable	X
Rack Related		
3812	Power Cable Group, Right, Drawer Position #4601	X
3802	Power Cable Group, Left, Drawer Position #4601	X
3811	Power Cable Group, Right, Drawer Position #4605	X

Feature code	Description	Status
3801	Power Cable Group, Left, Drawer Position #4605	X
3810	Power Cable Group, Right, Drawer Position #4609	X
3800	Power Cable Group, Left, Drawer Position #4609	X
3817	Power Cable Group, Right, Drawer Position #4613	X
3807	Power Cable Group, Left, Drawer Position #4613	X
3816	Power Cable Group, Right, Drawer Position #4631	X
3806	Power Cable Group, Left, Drawer Position #4631	X
3815	Power Cable Group, Right, Drawer Position #4627	X
3805	Power Cable Group, Left, Drawer Position #4627	X
3814	Power Cable Group, Right, Drawer Position #4623	X
3804	Power Cable Group, Left, Drawer Position #4623	X
3813	Power Cable Group, Right, Drawer Position #4619	X
3803	Power Cable Group, Left, Drawer Position #4619	X
3164	RIO to RIO-2 Remote I/O Cable, 1.0M	X
3146	RIO-2 (Remote I/O-2) Cable, 1.2M	X
3165	RIO to RIO-2 Remote I/O Cable, 3.0M	X
3168	RIO-2 (Remote I/O-2) Cable, 2.5M	X
SCSI disks		
3157	18.2 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W
3158	36.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W
3273	36.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly	X
3277	36.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly	X
3159	73.4 GB 10,000 RPM Ultra3 SCSI Disk Drive Assembly	W
3274	73.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly	X
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly	X
3275	146.8 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly	X

Feature code	Description	Status
Displays		
3628	IBM P260/P275 Color Monitor, Business Black, and Cable	X
3627	IBM P76/P77 Color Monitor, Business Black, Captured Cable	W
3637	IBM T541H 15" TFT Color Monitor, Business Black, Captured Cable	X
Memory		
4456	4 GB Memory Card	X
4457	8 GB Memory Card	X
4458	16 GB Memory Card	X
CEC related		
6199	DC Power Converter Assembly, CEC	X
Disk bays		
6569	ULTRA320 SCSI Backplane, two position	X
Software pre-install		
5005	Software Preinstall	X
Processors		
5513	4-Way 1.3GHz POWER4 Processor, 128 MB L3 Cache and System Planar	X
5518	4-Way 1.7GHz POWER4+ Processor, 128 MB L3 Cache and System Planar	X
5511	8-Way 1.1GHz POWER4 Processor, 128 MB L3 Cache and System Planar	X
5515	8-Way 1.5GHz POWER4+ Processor, 128 MB L3 Cache and System Planar	X
Rack related		
6591	Dual CEC side by side Mounting Enclosure	X
Specify codes		
9004	Southern Hemisphere Designator for Monitors	X

Feature code	Description	Status
Rack ID		
4601	Drawer Placement Indicator, 1U Position, Primary Rack	X
4605	Drawer Placement Indicator, 5U Position, Primary Rack	X
4609	Drawer Placement Indicator, 9U Position, Primary Rack	X
4613	Drawer Placement Indicator, 13U Position, Primary Rack	X
4619	Drawer Placement Indicator, 19U Position, Primary Rack	X
4623	Drawer Placement Indicator, 23U Position, Primary Rack	X
4627	Drawer Placement Indicator, 27U Position, Primary Rack	X
4631	Drawer Placement Indicator, 31U Position, Primary Rack	X
4651	Rack Indicator, Rack #1	X
4652	Rack Indicator, Rack #2	X
4653	Rack Indicator, Rack #3	X
4654	Rack Indicator, Rack #4	X
4655	Rack Indicator, Rack #5	X
4656	Rack Indicator, Rack #6	X
4657	Rack Indicator, Rack #7	X
4658	Rack Indicator, Rack #8	X
4659	Rack Indicator, Rack #9	X
4660	Rack Indicator, Rack #10	X
4661	Rack Indicator, Rack #11	X
4662	Rack Indicator, Rack #12	X
4663	Rack Indicator, Rack #13	X
4664	Rack Indicator, Rack #14	X
4665	Rack Indicator, Rack #15	X
4666	Rack Indicator, Rack #16	X

8.2 The RS/6000 SP in the IBM @server Cluster 1600

The RS/6000 SP system, participating in the IBM @server Cluster 1600 environment, delivers solutions to some of the most complex and large technical and commercial applications by simultaneously bringing dozens of RISC processing nodes to a computing problem. This parallel processing capability enhances computing performance and throughput many times over serial computing.

The basic SP building block is the processor node. Nodes consist of a processor, memory, disk, SP Switch2/SP Switch adapter slot(s), and PCI expansion slots for I/O and connectivity. Node types (Thin, Wide, or High) with varying processor power and I/O attachment options may be mixed within a system.

SP systems can be expanded and extended with RS/6000 and IBM @server pSeries through the application of the IBM @server Cluster 1600 environment (previously known as SP-Attached Servers). Supported servers include the machine type models 7017, 7026, 7038, 7039, 7028, and 7040. These servers can also be connected to the switch-only models of 9076-555/557 or 9076-556/558 for switched IBM @server Clusters. Scaling of servers is limited within the Cluster (9078) and SP (9076) environment.

8.2.1 The 9076 555/557 and 556/558 Models

The following are descriptions of the Models 555, 556, 557, and 558.

Model abstract 9076-555

The IBM 9076 Tall Frame Model 555 requires an SP Switch and a minimum of two clustered RS/6000 or pSeries servers along with a control workstation to create a functional system. An additional SP Switch is supported in the Model 555 expansion frame feature 1555. The 9076 Model 555 has a three-phase power system that includes N+1 power capability. This frame accepts power input 200-240V 50/60Hz 3-phase 50 A or 380-415V 50/60Hz 3-phase 30 A. Frame Redundant Power Input Features are available. SP nodes are supported on the SP Switch ports within the Model 555. SP Switches can connect to other SP Switches in other 9076 Models (for example, switch-to-switch connections between 9076 Model 555, and Model 550 and Model 557 are supported. Special attention to Cluster system planning for Model to Model cabling is required.)

Model abstract 9076-556

The IBM 9076 Tall Frame Model 556 requires an SP Switch2 and a minimum of two clustered RS/6000 or pSeries servers along with a control workstation to create a functional system. Additional SP Switch2s and features are supported for each Model 556 to scale the frame up to a maximum of eight SP Switch2s.

Both single plane and two-plane SP Switch2 environments are supported within the server scaling limits. The 9076 Model 556 has a three-phase power system that includes N+1 power capability. This frame accepts power input 200-240V 50/60Hz 3-phase 50 A or 380-415V 50/60Hz 3-phase 30 A. Frame Redundant Power Input Features are available. SP nodes are supported on the SP Switch2 ports within the Model 556. SP Switch2s can connect to other SP Switch2s in other 9076 Models (for example, switch-to-switch connections between 9076 Model 556, and Model 550 and Model 558 are supported. Special attention to Cluster system planning for Model to Model cabling is required.)

Model abstract 9076-557

The IBM 9076 Model 557 requires an SP Switch and a minimum of two clustered RS/6000 or pSeries servers along with a control workstation to create a functional system. The Model 557 is used for mounting the SP Switch and power sub-system in IBM's 19-in. racks, enhancing the available switch building blocks of the IBM @server Cluster 1600. This model is field installed into a new or existing 7014-T00 or 7014-T42 rack, requiring 16 EIA positions. The 9076-557 can have up to one SP Switch. The Model 557 requires two of the 7014 PDUs (power distribution units) for dual power input. The two power cords for Model 557 plug into one receptacle of each of the PDU units. SP nodes are supported on the SP Switch ports within the Model 557. The SP Switch within Model 557 can connect to other SP Switches in other 9076 Models (for example, switch-to-switch connections between 9076 Model 557, and Model 550 and Model 555 are supported). Special attention to Cluster system planning for Model to Model cabling is required.

Model abstract 9076-558

The IBM 9076 Model 558 requires an SP Switch2 and a minimum of two clustered RS/6000 or pSeries servers along with a control work station to create a functional system. The Model 558 is used for mounting the SP Switch2 and power sub-system in IBM's 19-in. racks, enhancing the available switch building blocks of the IBM @server Cluster 1600. This model is field installed into a new or existing 7014-T00 or 7014-T42 rack, requiring 16 EIA positions. The 9076-558 can have up to two SP Switch2s. The Model 558 requires two of the 7014 PDUs (power distribution units) for dual power input. The two power cords for Model 558 plug into one receptacle of each of the PDU units. Both single plane and two-plane SP Switch2 environments are supported within the server scaling limits. The SP Switch2s within Model 558 can connect to other SP Switch2s in other 9076 Models (for example, switch-to-switch connections between 9076 Model 558, and Model 550 and Model 556 are supported. Special attention to Cluster system planning for Model to Model cabling is required.)

SP Switch2 Adapter (FC 4012)

For models 550, 556, 558, and feature 1550.

This feature provides an SP Switch2 to support total expansion to 16 nodes within a single frame. Also supported are nodes in non-switched expansion frames, and IBM @server Cluster environments (formerly known as SP Attached Servers). SP Switch cables are separate features.

Limitations:

- ▶ Support SP PCI Thin, Wide, or High Nodes only.
- ▶ Cannot be mixed with other SP switch networks (features 4007, 4008, 4010, or 4011).
- ▶ The 4012 SPS2 switch is not supported within the 2.01 meter (79-in.) frames.

Requires AIX Version 4.3.3 with the 4330-02 Recommended Maintenance package or later, with IBM Parallel System Support Programs for AIX (PSSP) Version 3.2 plus IY12062 (select feature 9432). Support for Thin and Wide PCI nodes requires PSSP Version 3.4.

SP Switch MX2 Adapter (FC 4023)

This feature provides an SP Switch MX2 Adapter for installation in 375 MHz POWER3 SMP nodes (FC 2056, 2057, or 2058) or POWER3 SMP nodes (FC 2052, 2053, or 2054). This switch feature provides node connection to the SP Switch network. SP Switch Adapter (FC 4020), SP Switch MX Adapter (FC 4022), and SP Switch MX2 Adapter can co-exist in an SP configuration, along with SP Attached Servers.

This adapter requires AIX Version 4.3.2 with APAR IX85409 or later with PSSP Version 3.1 (select feature 9431).

SP Switch2 Adapter (FC 4025)

This feature provides an SP Switch2 Adapter for SP PCI High Nodes. This feature has a prerequisite of an SP Switch2 Interposer.

Limitations:

- ▶ Support SP Switch2 network only.
- ▶ Support a maximum of two SP Switch2 Adapters per SP PCI High Node in a two-plane configuration.

SP Switch2 MX2 Adapter (FC 4026)

This feature provides an SP Switch2 MX2 Adapter for installation in 375 MHz POWER3 SMP Thin and Wide nodes (FC 2056, 2057), POWER3 SMP Thin and

Wide nodes (FC 2052, 2053), or 332MHz SMP Thin and Wide nodes (FC 2050, 2051). This switch feature provides node connection to the SP Switch2 network. The SP Switch2 MX2 Adapter can co-exist in an SP configuration, along with the SP Switch2 Adapter and the SP Switch2 PCI Adapter in the SP Attached Servers.

This adapter requires PSSP Version 3.4. This feature has a prerequisite of an SP Switch2 Interposer.

Limitations:

- ▶ Support SP Switch2 network only.
- ▶ Support a maximum of one SP Switch2 MX2 Adapter per SP PCI Thin Node or Wide Node. This adapter does not support two-plane SP Switch2 configurations.

SP Switch2 Interposer (FC 4032)

This feature provides an interposer to connect each SP Switch2 Adapter to the SP Switch2. FC 4032 is also required to connect the SP Switch2 MX2 Adapter to the SP Switch2. FC 4032 is also required to connect the SP Switch2 PCI Adapter (within the SP Attached Servers) to the SP Switch2.

For an SP system with more than one SP Switch2 installed, 16 interposers are required for each SP Switch2 for the switch-to-switch connections.

8.3 7045-SW4 High-Performance Switch

The IBM @server pSeries High-Performance Switch (pSeries HPS) communication subsystem is an evolutionary step in network data transfers, using technology based on the proven architecture of SP Switch and SP Switch2. The pSeries HPS technology augments the IBM @server pSeries 690 and 655 clustered servers by increasing the communication bandwidth between servers and partitions within the cluster.

With the pSeries HPS, all network connections pass through an interface card packaged as either a 2-Link Switch Network Interface or a 4-Link Switch Network Interface (SNI). Each link on the SNI allows message passing between the server bus and the switch. These SNI packages are features of the p690 or p655 servers.

The pSeries HPS is delivered as a new machine type 7045, model SW4 that can be incorporated within the p690 system rack (7040-61R), the p655 system rack (7040-W42), or as a stand-alone multi-switch rack (7040-W42 without a p655 server.)

- ▶ The pSeries HPS (7045-SW4) is a 4 U tall, 24-in. drawer, rack-mounted device. It houses the High-Performance Switch planar and associated components. When the HPS is within a p690 rack or a p655 rack (7040-W42 with a 7039-651), the switch is installed with the bottom of the drawer at placement indicator (FC 4601), the 1 U position of either of the racks.
- ▶ When the HPS is within a 7040-W42 rack without a p655 (7039-651), up to eight HPSs can be installed in the primary rack. Up to 16 High-Performance Switches can be installed in the primary rack plus the Expansion Rack (FC 8691).

The HPS (7045-SW4) is added to the hardware models supported with the IBM @server pSeries Cluster 1600 running Cluster Systems Management (CSM) for AIX, Version 1.3.2

Planned availability date:

- ▶ October 31, 2003, for pSeries High-Performance Switch network supporting up to 16 servers, with up to 32 links. Cluster 1600 enhancements include the High-Performance Switch network in CSM clusters of p690 and p655 servers.
- ▶ July 30, 2004, for pSeries High-Performance Switch network supporting up to 64 servers (of which 32 servers can be p690), with up to 128 links. Also included are Frame Extender (FC 6234), Dual SNI to Single SNI Converter (FC 6437), and Expansion Rack (FC 8691) for the multi-switch rack (7040-W42).



Storage, communication, and other I/O adapters

This chapter describes the SCSI and SSA storage protocols and reports the storage, communications, and other I/O adapters that are currently available with the RS/6000 and pSeries systems.

Note: Refer to Appendix B “Adapter Placement Guidelines” or contact your local IBM representative for additional planning or hot-plug information.

Table 9-1 lists the available PCI I/O adapters sorted in numerical order.

Table 9-1 Available PCI I/O adapters and their feature codes

Feature code	PCI adapter name
2498	PCI 4-Channel Ultra3 SCSI RAID Adapter
2732	IBM Short-wave Serial HIPPI PCI Adapter for RS/6000
2751	S/390 ESCON Channel Adapter
2944	128-Port Asynchronous Adapter EIA-232
2946	Turboways 622 Mbps PCI MMF ATM Adapter

Feature code	PCI adapter name
2947	ARTIC960Hx 4-Port Selectable Adapter
2962	2-port Multiprotocol X.25 Adapter
2969	Gigabit Ethernet - SX
2975	10/100/1000 Base-T Ethernet PCI Adapter
4953	PCI ATM MMF Adapter 64-bit/66 MHz
4957	PCI ATM UTP Adapter 64-bit/66 MHz
4959	4/16 Mbps token-ring adapter
4960	E-business Cryptographic Accelerator
4961	Universal 4-Port 10/100 Ethernet adapter
4962	10/100 Mbps Ethernet PCI Adapter II
4963	PCI Cryptographic Coprocessor (FIPS-4)
5700	Gigabit Ethernet-SX PCI Adapter
5701	10/100/1000 Base-TX Ethernet PCI-X Adapter
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter
5707	2-Port Gigabit Ethernet-SX PCI-X Adapter
5709	Dual Channel SCSI RAID Enablement Card
5710	PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter
5711	PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter
5712	PCI-X Dual Channel Ultra320 SCSI Adapter
6203	PCI Dual Channel Ultra3 SCSI Adapter
6204	PCI Universal Differential Ultra SCSI
6205	PCI Dual Channel Ultra2 SCSI Adapter
6206	Ultra SCSI Single-Ended
6207	Ultra SCSI Differential
6228	Gigabit Fibre Channel Adapter for 64-bit PCI Bus (FC 6228)

Feature code	PCI adapter name
6230	Advanced SerialRAID Plus Adapter
6239	2 Gigabit Fibre Channel PCI-X Adapter
6310	ARTIC960RxD Quad Digital Trunk Adapter
6312	Quad Digital Trunk Telephony PCI Adapter
8396	IBM RS/6000 SP System Attachment Adapter
8397	SP Switch2 PCI Attachment Adapter
8398	SP Switch2 PCI-X Attachment Adapter

The PCI adapters are categorized into following groups:

- ▶ PCI Storage adapters
- ▶ System adapters
- ▶ Asynchronous adapters
- ▶ ARTIC adapters
- ▶ WAN adapters
- ▶ ATM adapters
- ▶ Ethernet and token-ring adapters
- ▶ Cryptographic adapters

9.1 PCI storage adapters

This section contains adapter descriptions of the PCI SCSI, SSA, and Fibre storage adapters for the RS/6000 and pSeries systems.

9.1.1 PCI-X Dual Channel Ultra320 SCSI RAID Adapter (FC 5703)

The PCI-X Dual Channel Ultra320 SCSI RAID Adapter (FC 5703) is a 64-bit 3.3 volt, bootable high performance Ultra320 SCSI RAID Adapter providing RAID 0, 5, or 10 capability, and can address up to thirty 16-bit SCSI physical disk drives on two independent SCSI buses.

To increase the data writing performance, a 40 MB non-volatile fast-write cache is provided as a resident part of this adapter. The 40 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 Dual Channel Ultra320 SCSI RAID Adapter has two independent Ultra320 SCSI buses. There are two internal ports and two external ports. The two internal ports are shared with the two external ports. The SCSI buses can drive either an internal port or an external port. The internal ports can be used to provide an internal RAID solution on certain supporting pSeries systems with internal multiple disk drives or packs of drives. Internally attached Ultra320 devices are designed to run at a data rate of up to 320 MBps on systems that have internal backplanes that are capable of supporting Ultra320 speeds.

In order to achieve an Ultra320 SCSI bus data rate of up to 320 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage Differential (LVD) drivers and receivers. To fully utilize this 320 MBps performance, all attaching devices should also be Ultra320 LVD devices. But, if Ultra2, Ultra3, or Ultra320 devices coexist on the same bus, each device will operate at its rated speed. For lower speed single-ended (SE) devices, the SCSI bus will switch to single-ended (SE) performance and interface to all devices on that SCSI bus at the lower SE bus data rate of the device.

When an array configuration is selected with the RAID Manager, the disk drives being designated as part of the array (attached to either the internal or external ports) are required to be formatted to 522 byte sectors. 522 byte sectors provide additional CRC error checking for improved data integrity. A menu option is provided in the AIX supporting software, which will reformat these disk drives prior to their usage in an array. Conversely, when a disk drive is removed from an

array, a similar menu option is also provided to re-format them back to 512 byte sectors.

Note: Some disk drives require that their microcode be updated to the latest level before being formatted to 522 Byte Sectors. Also, there are some disk drives that do not support being formatted to 522 Byte Sectors. The PCI-X SCSI Disk Array Manager will inform the user of these known situations when they exist. For disk microcode updates, go to the following URL:

<http://techsupport.services.ibm.com/server/mdownload/>

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 .3 meter converter cable, VHDCI to P, Mini-68 pin to 68-pin, (FC 2118) can be used with older external SE devices or subsystems to allow connection to the VHDCI connector on the PCI-X Dual Channel Ultra320 SCSI RAID Adapter.

The two external ports provide connectivity to an IBM 2104-DS4 Expandable Storage Plus Drawer or 2104-TS4 Expandable Storage Plus Tower at up to 320 MBps SCSI bus data rate configured as either a non-array or an array of disks. Also the two external ports provide non-array connectivity to numerous other SCSI external subsystems. Check the external subsystem sales/Web pages for verification of connectivity support with this adapter.

Limitations:

- ▶ The two external ports provide connectivity to an IBM 2104-DU3 Expandable Storage Plus Drawer or 2104-TU3 Expandable Storage Plus Tower at up to 160 MBps SCSI bus data rate as well as 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, but is limited to only non-array configuration support.
- ▶ The two external ports do not support the connection to the IBM 7131-105 IBM Multi-Storage Tower Model 105.
- ▶ Even though the Dual Channel Ultra320 SCSI RAID Adapter has ports that run at Ultra320 SCSI speeds (up to 320 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system disk backplane.
- ▶ Disk drives internal to the pSeries system shipped prior to September 1, 2003 require a disk drive microcode update to run at Ultra320 speed. To obtain the appropriate microcode update, go to the following URL:
<http://techsupport.services.ibm.com/server/mdownload/>
- ▶ Attributes provided: Attachment of internal and external SCSI devices.

- ▶ Attributes required: One PCI or PCI-X bus slot.

9.1.2 Dual Channel SCSI RAID Enablement Card (FC 5709)

The Dual Channel SCSI RAID Enablement Card (FC 5709) is a bootable high performance SCSI RAID Enablement feature providing RAID 0, 5, or 10 capability to select pSeries systems with the appropriate supporting integrated SCSI adapter and internal multiple disk drives or packs of drives.

To increase the data writing performance, a 16 MB non-volatile fast-write cache is provided as a resident part of this feature. The 16 MB fast-write cache can provide an 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.

When an array configuration is selected with the RAID Manager, the disk drives being designated as part of the array are required to be formatted to 522 byte sectors. 522 byte sectors provide additional CRC error checking for improved data integrity. A menu option is provided in the AIX supporting software, which will reformat these disk drives prior to their usage in an array. Conversely, when a disk drive is removed from an array, a similar menu option is also provided to re-format them back to 512 byte sectors.

Note: Some disk drives require that their microcode be updated to the latest level before being formatted to 522 byte sectors. Also, there are some disk drives that do not support being formatted to 522 byte sectors. The PCI-X SCSI Disk Array Manager will inform the user of these known situations when they exist. For disk microcode updates, go to the following URL:

<http://techsupport.services.ibm.com/server/mdownload/>

Limitations:

- ▶ Even though the supporting integrated adapter with the Dual Channel SCSI RAID Enablement Card has ports that run at Ultra320 SCSI speeds (up to 320 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system disk backplane.
- ▶ Disk drives internal to the pSeries system shipped prior to September 1, 2003 require a disk drive microcode update to run at Ultra320 speed. To obtain the appropriate microcode update, go to the following URL:
<http://techsupport.services.ibm.com/server/mdownload/>
- ▶ Attributes provided: Internal disk drives to be configured as an array.

- ▶ Attributes required: Systems with supporting integrated SCSI adapter.

9.1.3 PCI-X Dual Channel Ultra320 SCSI Blind Swap (FC 5710)

The PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter (FC 5710) is a 64-bit 3.3 volt adapter and is an excellent solution for high-performance SCSI applications. The PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter provides two SCSI channels (buses), each capable of running 320 MBps (max.), up to twice the maximum data transfer rate of the previous Dual Channel Ultra3 SCSI adapter (160 MBps). Each SCSI bus can either be internal (on systems that support internal SCSI device or backplane attachments) or external. Internally attached Ultra320 devices are designed to run at a data rate of up to 320 MBps on systems that have internal backplanes that are capable of supporting Ultra320 speeds.

In order to achieve an Ultra320 SCSI bus data rate of up to 320 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage Differential (LVD) drivers and receivers. To fully utilize this 320 MBps of performance, all attaching devices should also be Ultra320 LVD devices. But, if Ultra2, Ultra3, or Ultra320 devices coexist on the same bus, each device will operate at its rated speed. For lower speed single-ended (SE) devices, the SCSI bus will switch to single-ended (SE) performance and interface to all devices on that SCSI bus at the lower SE 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 .3 meter converter cable, VHDCI to P, Mini-68 pin to 68-pin, (FC 2118) can be used with older external SE devices or subsystems to allow connection to the VHDCI connector on the PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter.

The two external ports provide connectivity to an IBM 2104-DS4 Expandable Storage Plus Drawer or 2104-TS4 Expandable Storage Plus Tower at up to 320 MBps SCSI bus data rate. The two external ports also provide connectivity to an IBM 2104-DU3 Expandable Storage Plus Drawer or 2104-TU3 Expandable Storage Plus Tower at up to 160 MBps SCSI bus data rate and 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. Also, the two external ports provide connectivity to numerous other SCSI external subsystems. Check the external subsystem sales/Web pages for verification of connectivity support with this adapter.

The PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter (FC 5710) is a native boot adapter with AIX 5L Version 5.1 or AIX 5L Version 5.2 software in a supported pSeries or RS/6000 systems. The adapter also supports target mode.

Limitations:

- ▶ The two external ports do not support the connection to the IBM 7131-105 IBM Multi-Storage Tower Model 105.
- ▶ Even though the Dual Channel Ultra320 SCSI RAID Blind Swap Adapter has ports that run at Ultra320 SCSI speeds (up to 320 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system disk backplane.
- ▶ Disk drives internal to the pSeries system shipped prior to September 1, 2003 require a disk drive microcode update to run at Ultra320 speed. To obtain the appropriate microcode update, go to the following URL:
<http://techsupport.services.ibm.com/server/mdownload/>
- ▶ Attributes provided: Attachment of internal SCSI devices (on systems that support an internal SCSI device or backplane attachment with this adapter) and external SCSI devices.
- ▶ Attributes required: One available 3.3 volt PCI or PCI-X slot.

9.1.4 PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap (FC 5711)

The PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter (FC 5711) is a 64-bit 3.3 volt, bootable high performance Ultra320 SCSI RAID Adapter providing RAID 0, 5, or 10 capability and can address up to thirty 16-bit SCSI physical disk drives on two independent SCSI buses.

To increase the data writing performance, a 40 MB non-volatile fast-write cache is provided as a resident part of this adapter. The 40 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 Dual Channel Ultra320 SCSI RAID Blind Swap Adapter has two independent Ultra320 SCSI buses. There are two internal ports and two external ports. The two internal ports are shared with the two external ports. The SCSI buses can drive either an internal port or an external port. The internal ports can be used to provide an internal RAID solution on certain supporting pSeries systems with internal multiple disk drives or packs of drives. Internally attached Ultra320 devices are designed to run at a data rate of up to 320 MBps on systems that have internal backplanes that are capable of supporting Ultra320 speeds.

In order to achieve an Ultra320 SCSI bus data rate of up to 320 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage

Differential (LVD) drivers and receivers. To fully utilize this 320 MBps performance, all attaching devices should also be Ultra320 LVD devices. But, if Ultra2, Ultra3, or Ultra320 devices coexist on the same bus, each device will operate at its rated speed. For lower speed single-ended (SE) devices, the SCSI bus will switch to single-ended (SE) performance and interface to all devices on that SCSI bus at the lower SE bus data rate of the device.

When an array configuration is selected with the RAID Manager, the disk drives being designated as part of the array (attached to either the internal or external ports) are required to be formatted to 522 byte sectors. 522 byte sectors provide additional CRC error checking for improved data integrity. A menu option is provided in the AIX supporting software, which will reformat these disk drives prior to their usage in an array. Conversely, when a disk drive is removed from an array, a similar menu option is also provided to re-format them back to 512 byte sectors.

Note: Some disk drives require that their microcode be updated to the latest level before being formatted to 522 byte sectors. Also, there are some disk drives that do not support being formatted to 522 byte sectors. The PCI-X SCSI Disk Array Manager will inform the user of these known situations when they exist. For disk microcode updates, go to the following URL:

<http://techsupport.services.ibm.com/server/mdownload/>

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 .3 meter converter cable, VHDCI to P, Mini-68 pin to 68-pin, (FC 2118) can be used with older external SE devices or subsystems to allow connection to the VHDCI connector on the PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter.

The two external ports provide connectivity to an IBM 2104-DS4 Expandable Storage Plus Drawer or 2104-TS4 Expandable Storage Plus Tower at up to 320 MBps SCSI bus data rate configured as either a non-array or an array of disks. Also, the two external ports provide non-array connectivity to numerous other SCSI external subsystems. Check the external subsystem sales/Web pages for verification of connectivity support with this adapter.

Limitations:

- ▶ The two external ports provide connectivity to an IBM 2104-DU3 Expandable Storage Plus Drawer or 2104-TU3 Expandable Storage Plus Tower at up to 160 MBps SCSI bus data rate, as well as 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, but is limited to only non-array configuration support.

- ▶ The two external ports do not support the connection to the IBM 7131-105 IBM Multi-Storage Tower Model 105.
- ▶ Even though the Dual Channel Ultra320 SCSI RAID Blind Swap Adapter has ports that run at Ultra320 SCSI speeds (up to 320 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system disk backplane.
- ▶ Disk drives internal to the pSeries system shipped prior to September 1, 2003 require a disk drive microcode update to run at Ultra320 speed. To obtain the appropriate microcode update, go to the following URL:
<http://techsupport.services.ibm.com/server/mdownload/>
- ▶ Attributes provided: Attachment of internal and external SCSI devices.
- ▶ Attributes required: One PCI or PCI-X bus slot.

9.1.5 PCI-X Dual Channel Ultra320 SCSI Adapter (FC 5712)

The PCI-X Dual Channel Ultra320 SCSI Adapter (FC 5712) is a 64-bit 3.3 volt adapter and is an excellent solution for high-performance SCSI applications. The PCI-X Dual Channel Ultra320 SCSI Adapter provides two SCSI channels (buses), each capable of running 320 MBps (max.), up to twice the maximum data transfer rate of the previous Dual Channel Ultra3 SCSI adapter (160 MBps). Each SCSI bus can either be internal (on systems that support internal SCSI device or backplane attachments) or external. Internally attached Ultra320 devices are designed to run at a data rate of up to 320 MB per second on systems that have internal backplanes that are capable of supporting Ultra320 speeds.

In order to achieve an Ultra320 SCSI bus data rate of up to 320 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage Differential (LVD) drivers and receivers. To fully utilize this 320 MBps performance, all attaching devices should also be Ultra320 LVD devices. But, if Ultra2, Ultra3, or Ultra320 devices coexist on the same bus, each device will operate at its rated speed. For lower speed single-ended (SE) devices, the SCSI bus will switch to single-ended (SE) performance and interface to all devices on that SCSI bus at the lower SE 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 .3 meter converter cable, VHDCI to P, Mini-68 pin to 68-pin, (FC 2118) can be used with older external SE devices or subsystems to allow connection to the VHDCI connector on the PCI-X Dual Channel Ultra320 SCSI Adapter.

The two external ports provide connectivity to an IBM 2104-DS4 Expandable Storage Plus Drawer or 2104-TS4 Expandable Storage Plus Tower at up to 320

MBps SCSI bus data rate. The two external ports also provide connectivity to an IBM 2104-DU3 Expandable Storage Plus Drawer or 2104-TU3 Expandable Storage Plus Tower at up to 160 MBps SCSI bus data rate and 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. Also, the two external ports provide connectivity to numerous other SCSI external subsystems. Check the external subsystem sales/Web pages for verification of connectivity support with this adapter.

The PCI-X Dual Channel Ultra320 SCSI Adapter (FC 5712) is a native boot adapter with AIX 5L Version 5.1 or AIX 5L Version 5.2 software in a supported pSeries or RS/6000 systems. The adapter also supports target mode.

Limitations:

- ▶ The two external ports do not support the connection to the IBM 7131-105 IBM Multi-Storage Tower Model 105.
- ▶ Even though the Dual Channel Ultra320 SCSI RAID Adapter has ports that run at Ultra320 SCSI speeds (up to 320 MBps), the internally attached disk drives will run at a maximum SCSI bus data rate specified by that supporting system disk backplane.
- ▶ Disk drives internal to the pSeries system shipped prior to September 1, 2003 require a disk drive microcode update to run at Ultra320 speed. To obtain the appropriate microcode update, go to the following URL:
<http://techsupport.services.ibm.com/server/mdownload/>
- ▶ Attributes provided: Attachment of internal SCSI devices (on systems that support an internal SCSI device or backplane attachment with this adapter) and external SCSI devices.
- ▶ Attributes required: One available 3.3 volt PCI or PCI-X slot.

9.1.6 PCI 4-Channel Ultra3 SCSI RAID Adapter (FC 2498)

The RS/6000 PCI 4-Channel Ultra3 SCSI RAID Adapter (FC 2498) is a non-bootable high performance Ultra3 SCSI RAID Adapter providing RAID 0, 1, 1E, 5, or 5E capability, and can address up to sixty 16-bit SCSI physical disk drives on four independent SCSI buses.

To increase the data writing performance, a 128 MB fast-write cache is provided as a resident part of this adapter. The 128 MB fast-write cache is a resident feature of the PCI 4-Channel Ultra3 SCSI RAID Adapter that utilizes non-volatile RAM. During the unlikely event of an PCI 4-Channel Ultra3 SCSI RAID Adapter failure, a replacement PCI 4-Channel Ultra3 SCSI RAID Adapter can be installed and the fast-write cache can be removed from the failing adapter and installed in the new adapter, ensuring data integrity. The 128 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 4-Channel Ultra3 SCSI RAID Adapter has four independent Ultra3 SCSI buses. There are two internal ports and four external ports. The two internal ports are shared with two of the external ports. Two of the four buses can drive either an internal port or an external port. The other two buses only drive external ports. The internal ports can be used to provide an internal RAID solution on supporting RS/6000 systems. Systems with one or two internal 6-pack disks can attach to a PCI 4-Channel Ultra3 SCSI RAID Adapter. The four external ports provide connectivity to an IBM 2104-DU3 Expandable Storage Plus Drawer or 2104-TU3 Expandable Storage Plus Tower at up to 160 MBps SCSI bus data rate. The four external ports also 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.

Configuration notes

- ▶ The four external ports do not support the connection to the IBM 7131-105 external Fast/Wide SCSI disk enclosure.
- ▶ Even though the PCI 4-Channel Ultra3 SCSI RAID Adapter has ports that run at Ultra3 SCSI speeds (up to 160 MBps) and 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.

Note: It is also possible to upgrade an old PCI 3-Channel Ultra SCSI RAID Adapter (FC 2494) to the newer PCI 3-Channel Ultra2 SCSI RAID Adapter (FC 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.

Hardware requirements

One free PCI slot is required for this feature.

Software requirements

AIX Version 4.3.3 or AIX 5L Version 5.1 or later is required for this feature.

Note: To have full support, and to take full advantage of the Ultra3 (up to 160 MBps) speed of this adapter, the proper AIX level also needs to be considered. AIX Version 4.3.3 with appropriate APAR updates or later AIX versions support the full range of SCSI bus data rates (including Ultra3 SCSI up to 160 MBps).

9.1.7 PCI Dual Channel Ultra3 SCSI Adapter (FC 6203)

The PCI Universal Differential Ultra SCSI Adapter is the latest technology advancement of an RS/6000 SCSI-2 differential adapter with a maximum data transfer rate of 40 MBps. This adapter has the capability to be plugged into the newer +3.3 volt PCI slots as well as the older +5 volt PCI slots. FC 6204 allows connection up to 15 external differential SCSI-2 F/W or Ultra SCSI type devices up to 25 meters away. FC 6204 will negotiate with each external device and transfer data at the fastest SCSI data transfer rate capable by the external device. Check the system sales pages as to which external differential subsystems are supported.

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).

Hardware requirements

One free PCI slot is required for this feature.

Software requirements

AIX 5L Version 5.1 or later is required for this feature.

9.1.8 PCI Universal Differential Ultra SCSI Adapter (FC 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 (FC 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).

Feature characteristics

The following are the major characteristics of this feature:

- ▶ 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

Software requirements

The following are the software requirements of this feature:

- ▶ AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.1.1 or later on the SP nodes

9.1.9 PCI Dual Channel Ultra2 SCSI Adapter (FC 6205)

The PCI Dual Channel Ultra2 SCSI Adapter (FC 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 (buses). 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 the attachment of various LVD and SE external subsystems. A three meter converter cable, VHDCI to P, Mini-68 pin to 68 pin, (FC 2118) can be used with older external SE subsystems to allow connection to the VHDCI connector on the PCI Dual Channel Ultra2 SCSI Adapter.

Note: To connect a FC 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 (FC 9799) and one of the appropriate cables: 9137 1.0 m Ultra2 SCSI cable or 9149 1.0 m SCSI-2 F/W cable.

Hardware requirements

One available PCI slot is required for this feature.

Software requirements

AIX Version 4.3.3 or AIX 5L Version 5.1 or later is required for this feature.

Notes: Any supported RS/6000 system can be set up to boot from the PCI Dual Channel Ultra2 SCSI Adapter (FC 6205).

- ▶ If you are running with AIX Version 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 Version 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 into 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.

9.1.10 PCI Single-Ended Ultra SCSI Adapter (FC 6206)

The PCI Single-Ended Ultra SCSI Adapter (FC 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.

Feature characteristics

The following are the key characteristics of this feature:

- ▶ 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

Configuration notes

The following are the key configuration notes for this feature:

- ▶ 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.

Cable options

The following optional cables are available for the (FC 6206) SCSI adapter:

- ▶ (FC 2117) 16-bit SE external Y-cable, 0.9 m
- ▶ (FC 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m
- ▶ (FC 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.

Hardware requirements

One available PCI slot is required for this feature.

Software requirements

AIX Version 4.3.3 or AIX 5L Version 5.1 or later is required for this feature.

9.1.11 PCI Differential Ultra SCSI Adapter (FC 6207)

The PCI Differential Ultra SCSI Adapter (FC 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.

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

Configuration notes

- ▶ Data transfer rates with FC 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.

Cable options

The following optional cables are available for the PCI SCSI-2 Ultra/Wide Differential Adapter:

- ▶ (FC 2114) 16-bit DE external Y-cable, 0.9 m
- ▶ (FC 2424) 16-bit adapter-to-adapter SCSI cable, 0.6 m long
- ▶ (FC 2425) 16-bit adapter-to-adapter SCSI cable, 2.5 m long

All cables conform to X3T9.2/90-048 standards.

Hardware requirements

One available PVCI slot

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later

9.1.12 Advanced SerialRAID Plus Adapter (FC 6230)

The Advanced SerialRAID Plus Adapter (FC 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 (FC 6235) for dual or single attach configurations
- ▶ Optional 128 MB DRAM (FC 6231) (standard 64 MB DRAM) for full use of 32 MB cache in 2-way operation

In addition, these functions were already supported on the FC 6225 and are now also supported on the FC 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 using 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 is not 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 second 70/30 Random 4 Kb I/Os/second, 50 MBps Write data rate, 80 MBps Read data rate.
- ▶ 2-way: 7500 I/Os per second 70/30 Random 4 Kb I/Os/second, 65 MBps Write data rate, 150 MBps Read data rate.

Note: These are lab measurements and may not be realized in all customer environments.

The Advanced SerialRAID Plus Adapter FC 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 FC 6230 can coexist on the same loop with the Advanced SerialRAID adapter FC 6225, as well as with the SSA RAID EL adapters FC 6215 and FC 6219. Note that 2-way Fast-Write Cache and many other capabilities are not supported with FC 6215 and FC 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 (FC 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 (FC 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.

Hardware requirements

One available PCI slot

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later

Note: Any supported RS/6000 system can be set up to boot from an Advanced SerialRAID Plus Adapter (FC 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 Version 4.3.3 or later, software native boot capability is supported.
- ▶ If your system is running AIX Version 4.2.1 or AIX Version 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.

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 9-2 shows an overview of the compatibility and use of both Optical Extender generations.

Table 9-2 Overview of the compatibility and use of both Optical Extenders

Optical Extender Type	Mode	SSA Fibre Optical Extender	SSA Advance Fibre Optical Extender
Feature Code		FC 5500	FC 8851
Link Speed	20 MBps	Yes	Yes
	40 MBps	Not supported	Yes (only with Advanced SerialRAID Adapter FC 6225 and 7133 Model D40/T40)
Fibre Cable Type	multimode	Either 50 μm or 62.5 μm	Either 50 μm or 62.5 μm (Mode Conditioning Patch Cords FC 8852 or FC 8853 must be installed)
	singlemode	Not supported	9 μm
Max. Distance	multimode	Up to 2.4 km	Up to 3 km with FC 6225 or FC 6230; up to 2.4 km with all other adapter types
	singlemode	Not supported	Up to 10 km (only with Advanced SerialRAID adapter FC 6225 or FC 6230, and 7133 Model D40/T40)
Max. Overall Attenuation		3 dB	8 dB
Supported drawers		All models	All models
Supported Adapter Types		All	All

9.1.13 Gigabit Fibre Channel Adapter for 64-bit PCI Bus (FC 6228)

The Gigabit Fibre Channel Adapter for 64-bit PCI Bus is a 64-bit address/data, short form factor PCI adapter with LC type external fiber connectors that provides single or dual initiator capability over an optical fiber link or loop running up to 100 Mbps. With the use of appropriate optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage. Distances of up to 500 meters are supported. When used with IBM supported Fibre Channel Storage Hub and switches supporting longwave optics, distances of up to 10 kilometers are capable.

The Gigabit Fibre Channel for 64-bit PCI Bus can be used to attach devices either directly, or by means of Fibre Channel Switches. If attaching a device or switch with a SC type fiber connector(s), use of an LC-SC Fibre Channel Conversion Cable (FC 2456) is required.

Refer to the following IBM storage subsystem Web page for additional supported server attachment information for IBM devices.

<http://www.storage.ibm.com/hardsoft/disk/products.htm>

Hardware requirements

One available PCI slot

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later

9.2 System adapters

The following is the description of the available system adapters.

9.2.1 IBM Short-wave Serial HIPPI PCI Adapter (FC 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 (FC 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.

Software requirements

- ▶ AIX Version 4.3.2 or later
- ▶ IBM HIPPI PCI Connectivity Version 4.3 AIX LPP (5765-E07)

9.2.2 S/390 ESCON Channel Adapter (FC 2751)

The PCI S/390 ESCON Channel Adapter (FC 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.

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

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

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 Version 2 (5765-D49) for mainframe connectivity
 - AIX program feature: PCI ESCON Tape Attachment Version 2 (5765-E04) for tape subsystem attachment

Software requirements

- ▶ Supported with AIX Version 4.2.1 and above on selected hardware platforms
- ▶ One of the following:
 - PCI ESCON Control Unit Connectivity Version 2 (separately ordered as LPP 5765-D49)
 - PCI ESCON Tape Attachment Version 2 (separately ordered as LPP 5765-D49) for tape subsystem attachment

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.

9.2.3 IBM RS/6000 SP System Attachment Adapter (FC 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).

Only one SP System Attachment Adapter (FC 8396) is allowed per RS/6000 system.

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

Software requirements

- ▶ AIX Version 4.3 or AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.1 or later

9.2.4 SP Switch2 PCI Attachment Adapter (FC 8397)

This adapter allows attachment to the SP Switch2. It allows the system to be integrated with an SP system and act as an attached SMP server within the SP infrastructure. It also includes diagnostics and publications.

Hardware requirements

- ▶ One free PCI slot
- ▶ SP Switch2
- ▶ Interconnection cables
- ▶ One Ethernet Adapter
- ▶ IBM Parallel System support programs for AIX.

Software requirements

AIX 5L Version 5.1 or later is required.

9.2.5 SP Switch2 PCI-X Attachment Adapter (FC 8398)

This adapter allows attachment to the SP Switch2 through the PCI-X bus. It also includes diagnostics and publications.

Hardware requirements

- ▶ One PCI-X slot
- ▶ SP Switch2 fabric on the Cluster

- ▶ PSSP (IBM Parallel System Support Programs for AIX.)

Software requirements

- ▶ AIX 5L Version 5.1 or AIX 5L Version 5.2 or later
- ▶ PSSP Version 3.4 or PSSP Version 3.5 or later

Notes:

- ▶ The 7028-6C4 with FC 8398 requires PSSP Version 3.4 or PSSP Version 3.5 with the following APARs for May 2003:
 - IY42367 for PSSP Version 3.4 support for p630 with SP Switch2 PCI-X in RIO mode
 - IY42368 for PSSP Version 3.5 support for p630 with SP Switch2 PCI-X in RIO mode.
- ▶ The 7028-6C4 with FC 8398 requires PSSP Version 3.4 or PSSP Version 3.5 with the following APARs for July 2003:
 - IY42359 PSSP Version 3.4 support for p630 with SP Switch2 PCI-X in RIO-2 mode
 - IY42358 PSSP Version 3.5 support for p630 with SP Switch2 PCI-X in RIO-2 mode

9.3 Asynchronous adapters

The following sections discuss the available asynchronous adapters.

9.3.1 8-Port Asynchronous Adapter EIA 232/RS-422 (FC 2943)

The 8-port Asynchronous Adapter (FC 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 that 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.

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

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

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.

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.4 or later on SP nodes
- ▶ Adapter device driver LPP image (provided with adapter)

9.3.2 128-Port Asynchronous Adapter EIA-232 (FC 2944)

The 128-port asynchronous adapter (FC 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.

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

Customer-supplied components

Feature (FC 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 9-3 and Table 9-4 provide the related feature codes and their descriptions.

Table 9-3 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.) ^a
8133	RJ-45 to DB-25 Converter Cable
8136	Rack Mountable Remote Asynchronous Node 16-Port EIA-232

a. This cable may be substituted for the 4.5 meter (15 foot) asynchronous controller cable whenever a customer configuration requires stacked Remote Asynchronous Nodes.

Table 9-4 The 2.4 Mbps RANs and cables

Feature Code	Description
8137	Enhanced Remote Asynchronous Node 16-Port EIA-232

Feature Code	Description
8138	Enhanced Remote Asynchronous Node 16-Port RS-422
2934	Asynchronous Terminal/Printer Cable, EIA-232
3926	Asynchronous 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

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.4 or later on SP Nodes
- ▶ Adapter device driver LPP image (provided with adapter)

9.4 ARTIC adapters

The following sections discuss the available ARTIC adapters.

9.4.1 ARTIC960Hx 4-Port Selectable Adapter (FC 2947)

The ARTIC960Hx 4-Port Selectable PCI Adapter (FC 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.

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

Customer-supplied cables

A connecting cable is required. The cables listed in Table 9-5 are available from IBM.

Table 9-5 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

Software requirements

- ▶ AIX Version 4.3.2 and APAR IX81860 or later for SDLC or Bisync
- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later
- ▶ 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.

9.4.2 ARTIC960RxD Quad Digital Trunk Adapter (FC 6310)

The ARTIC960RxD Quad Digital Trunk Adapter (FC 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.

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

Feature components

- ▶ One ARTIC960RxD Adapter (FC 6310).
- ▶ A connecting cable (as listed in Table 9-6) is required.

Table 9-6 *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
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

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later
- ▶ DirectTalk for AIX Version 2.1 LPP (5765-B81) or later to provide voice processing
- ▶ Adapter device driver (provided with adapter)

Quad Digital Trunk Telephony PCI Adapter (FC 6312)

The Quad Digital Trunk Telephony PCI Adapter is a highly integrated, intelligent IO adapter designed for use in computer telephony applications. The adapter is a 4 port, full length, universal PCI 2.2 compliant adapter. It performs voice processing for up to four T1 or E1 digital trunks, providing connectivity for 96 (T1) or 120 (E1) voice channels in a single PCI slot. The adapter is made up of two separate cards. A base card that interfaces with the host system and performs telephony processing functions, and a daughter card that provides the physical interface to the switch. The voice processing function is provided by WebSphere® Voice Response for AIX LPP with Direct Talk Technology. In conjunction with this adapter, network attachment cables using industry standard RJ-48 connector can be obtained from commercial cable suppliers in a variety of lengths to suit the particular installation.

Note: When two to four Quad Digital Trunk Telephony PCI Adapters (FC 6312) are ordered and planned to be used in a single partition, the Artic960RxD Quad DTA, H.100, 4-drop Cable (FC 2877) is required. When more than four Quad Digital Trunk Telephony PCI Adapters (FC 6312) are ordered and planned to be used in a single partition, the H.100 Bus 8-Position Cable (FC 4353) is required.

Configuration notes

- ▶ The Quad Digital Trunk Telephony PCI Adapter (FC 6312) cannot reside in the same system with the IBM ARTIC960RxD Quad Digital Trunk PCI Adapter (FC 6310).
- ▶ The Quad Digital Trunk Telephony PCI Adapter (FC 6312) is only supported in machines running in a full system partition.

Hardware requirements

- ▶ One available PCI slot

Software requirements

AIX Version 4.3.3 or AIX 5L Version 5.1 or later

- ▶ WebSphere Voice Response for AIX LPP application software.

9.5 WAN adapters

The following sections discuss the available WAN adapters.

9.5.1 2-Port Multiprotocol X.25 Adapter (FC 2962)

The 2-Port Multiprotocol Adapter (FC 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.

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

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 9-7.

Table 9-7 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)

Software requirements

- ▶ AIX Version 4.3.3 or later
- ▶ PSSP Version 3.4 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)

9.6 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) affects the class of networking services provided by the network.

The following is a list of the currently supported ATM adapters:

- ▶ Turboways 622 Mbps PCI MMF ATM Adapter (FC 2946)

- ▶ Turboways 155 Mbps PCI MMF ATM Adapter 64-bit/66 MHz (FC 4953)
- ▶ Turboways 155 Mbps PCI UTP ATM Adapter 64-bit/66 MHz (FC 4957)

For additional information, refer to *RS/6000 ATM Cookbook*, SG24-5525.

9.6.1 Turboways 622 Mbps PCI MMF ATM Adapter (FC 2946)

The IBM Turboways 622 Mbps PCI MMF ATM Adapter is a 64-bit, Universal PCI Adapter. This adapter provides direct access to the ATM network at a dedicated 622 Mbps full-duplex connection. The Turboways 622 Mbps PCI MMF ATM Adapter is a short form-factor adapter that interfaces to the system using the PCI bus and connects to the 622 Mbps ATM network using dual-SC type, multi-mode fiber cables. The Turboways 622 Mbps PCI MMF ATM Adapter utilizes 16 MB of SDRAM for control and 16 MB of SDRAM for packet memory. This ATM adapter also provides a hardware assist for TCP checksum, which can provide a performance improvement by minimizing the host CPU cycles.

Software requirements

- ▶ AIX Version 4.3.3 with 4330-05 Recommended Maintenance Package or AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.1 or later on SP nodes

9.6.2 IBM 64-bit/66 MHz PCI ATM 155 UTP Adapter (FC 4953)

The IBM 64-bit/66MHz PCI ATM 155 UTP Adapter provides dedicated, 155 Mbps full-duplex connection to ATM networks over either permanent virtual circuits (PVC) or switched virtual circuits (SVC). This adapter enables TCP/IP to run over an ATM network with Category-5 Unshielded Twisted Pair (UTP). It also supports communication with devices located on an ATM network or bridged to a Token-Ring, Ethernet, or other LAN.

Hardware requirements

One PCI slot

Software requirements

AIX 5L Version 5.1 or later

9.6.3 IBM 64-bit/66 MHz PCI ATM 155 MMF Adapter (FC 4957)

The IBM 64-bit/66MHz PCI ATM 155 MMF Adapter provides direct access to ATM networks. This 155 Mbps PCI ATM MMF Adapter provides dedicated 155 Mbps full-duplex connection using permanent virtual circuits (PVC) or switched

virtual circuits (SVC) and enables TCP/IP to run over an ATM network. The adapter also supports communication with devices located on an ATM network or bridged to a Token-Ring, Ethernet, or other LAN.

This 155 Mbps PCI MMF ATM Adapter is compatible with:

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

Hardware requirements

- ▶ One PCI slot

Software requirements

- ▶ AIX 5L Version 5.1 or later

9.7 Ethernet and token-ring adapters

The following section discusses the available Ethernet and token-ring adapters.

9.7.1 Gigabit Ethernet - SX Adapter (FC 2969)

The PCI Gigabit Ethernet - SX Adapter (FC 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.

Feature characteristics

- ▶ Compatible with IEEE 802.3z standards
- ▶ Supports full duplex operation over 1000BaseSX networks
- ▶ Supports jumbo frames with AIX Version 4.3.2 device driver or later
- ▶ Works in 32- or 64-bit slots

Feature components

- ▶ Adapter card
- ▶ Fiber wrap plug

- ▶ Installation instructions

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

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later
- ▶ PSSP Version 3.4 or later on SP nodes

9.7.2 10/100/1000 Base-T Ethernet PCI Adapter (FC 2975)

10/100/1000 Base-T Ethernet PCI Adapter is a Full Duplex Gigabit Ethernet adapter designed with highly integrated components to optimize cost and performance. The adapter interfaces to the system using the PCI bus and connects to the network using a 4-pair CAT-5 Unshielded Twisted Pair (UTP) cable for distances of up to 100 m. The 10/100/1000 Base-T Ethernet PCI Adapter supports jumbo frames for full duplex Fast and Gigabit Ethernet.

Note: For optimum performance, adapter should be placed in a 64-bit PCI card slot.

Limitation: AIX's Network Install Manager (NIM) boot is not supported with this adapter. The 1000 Mbps speed is not supported in Half Duplex (HDX) mode.

Hardware requirements

- ▶ One free PCI slot

Software requirements

- ▶ AIX Version 4.3.3 with 4330-05 Recommended Maintenance Package or AIX 5L Version 5.1

9.7.3 4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951)

The IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter for RS/6000 (FC 4951) makes four Ethernet ports available using a single PCI slot. This adapter provides the same function as purchasing four Ethernet adapters (FC 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.

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

Feature components

- ▶ Adapter card
- ▶ Device driver for adapter card
- ▶ Firmware update
- ▶ Installations instructions

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

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/>

9.7.4 4/16 Mbps Token-Ring Adapter (FC 4959)

The IBM token-ring PCI adapter for RS/6000 (FC 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.

Feature characteristics

- ▶ Communicates at 4 or 16 Mbps
- ▶ Supports both UTP-5 (RJ45) and STP (DB9) cable connections
- ▶ Auto-detects 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

Software requirements

- ▶ AIX Version 4.3.3 with 4330-08 Recommended Maintenance Package or later, or AIX 5L Version 5.1 or later.

9.7.5 IBM 4-Port 10/100 Ethernet Adapter (FC 4951)

IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter for RS/6000 is a single slot, long, 64-bit 33 MHz PCI adapter supporting four industry standard Ethernet 10Base-T or 100 Base-T interfaces either half or full duplex. Each port is provided with its own 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.

Network boot capability and Network Install Manager (NIM) capability are available using this adapter if no specific limitation is stated.

The IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951) should be considered where maximum port density is required per I/O card slot. But, for high end systems, where card slots are not the limiting factor and maximum throughput is required, the single port IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962) is the preferred solution.

Below are some performance factors to consider when choosing the right adapter for your needs. These performance comparisons are based on all four ports of the IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951) being active.

- ▶ With all four ports under full bandwidth conditions, a single IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951) is expected to deliver up to three times the comparable performance of a single IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962).
- ▶ Under most conditions, each port of the IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter (FC 4951) is expected to perform at greater than 50% of the throughput of the single port of the IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962).

Note: The resulting performance in your environment compared to the above may vary and depends upon the RS/6000 model, the I/O configuration, and associated workload of your applications.

Hardware requirements

- ▶ One free PCI slot

Software requirements

- ▶ AIX Version 4.3.3 or AIX 5L Version 5.1 or later

9.7.6 IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961)

IBM Universal 4-Port 10/100 Ethernet Adapter is a single slot, long, 64-bit, 33 MHz PCI adapter supporting four industry standard Ethernet 10 Base-T or 100 Base-T interfaces either half or full-duplex. Each port is provided with its own 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. Network boot capability and Network Install Manager (NIM) capability are available using this adapter if no specific limitation is stated. The IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961) should be considered where maximum port density is required per I/O card slot. But, for high end systems, where card slots are not the limiting factor and maximum throughput is required, the single port IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962) is the preferred solution.

Below are some performance factors to consider when choosing the right adapter for your needs. These performance comparisons are based on all four ports of the IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961) being active.

- ▶ With all four ports under full bandwidth conditions, a single IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961) is expected to deliver up to three times the comparable performance of a single IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962).
- ▶ Under most conditions, each port of the IBM Universal 4-Port 10/100 Ethernet Adapter (FC 4961) is expected to perform at greater than 50% of the throughput of the single port of the IBM 10/100 Mbps Ethernet PCI Adapter II (FC 4962).

Note: The resulting performance in your environment compared to the above may vary and depends upon the RS/6000 model, the I/O configuration, and associated workload of your applications.

Hardware requirements

- ▶ One free PCI slot

Software requirements

- ▶ AIX 5L Version 5.1 or later

9.7.7 10/100 Mbps Ethernet PCI Adapter II (FC 4962)

The 10/100 Mbps Ethernet PCI Adapter II is a small form factor, single port PCI Ethernet adapter. This high performance, low power Ethernet 10/100 Mbps LAN adapter can be used in both client and server PCI systems. The 10/100 Mbps Ethernet PCI Adapter II provides both 10Base-T and 100Base-TX full-duplex Ethernet LAN connectivity. The adapter supports Category-5 unshielded twisted pair cabling for both 10/100 Mbps and Category-3 unshielded twisted pair cabling for 10 Mbps. The 10/100 Mbps Ethernet PCI Adapter II supports:

- ▶ Half-/full-duplex 10/100 Mbps Ethernet interface
- ▶ 10/100 Mbps data rates
- ▶ Auto-negotiation for 10/100 speed and half/full-duplex
- ▶ Network boot capability and Network Install Manager (NIM)
- ▶ IEEE 802.3 Ethernet Specification
- ▶ IEEE 802.3u Fast Ethernet Specification
- ▶ After April 26, 2002, the 10/100 Mbps Ethernet PCI Adapter II supports the off-load of IP Security cryptographic algorithms by providing hardware assistance in performing data encryption and authentication. This support is provided with AIX 5L Version 5.1 (with appropriate software updates) and later. This IP Security function, normally performed with encryption software by the host, is off-loaded to this adapter to enhance network traffic throughput and reduce CPU utilization. If you are running with AIX 5L Version 5.1, to invoke the IP Security function on the adapter, you must obtain AIX 5L Version 5.1 software updates IY27069 and IY26514 or the 5100-02 Recommended Maintenance package. These updates can be obtained by ordering APAR IY28102, or by ordering the AIX 5L Version 5.1 Update CD (LCD4-1103-03) dated April 2002 or later.

Note: This IP Security function is not supported with AIX Version 4.3 software.

Hardware requirements

- ▶ One free PCI slot

Software requirements

- ▶ AIX 5L Version 5.1 or later
- ▶ An additional function called 'Large Send', sometimes known as TCP Segmentation, is also available. This function offloads the TCP segmentation operation from the AIX IP layer to the adapter for outgoing (transmit side) TCP segments. Another function known as checksum offload, which offloads the TCP/UDP Checksum Operation or workload from the CPU to the adapter,

is also available. Both of these functions are available with AIX 5L Version 5.1 with APAR IY38248 or later software, or AIX 5L Version 5.2 with APAR IY38492 or later software.

9.7.8 IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700)

The IBM Gigabit Ethernet-SX PCI-X Adapter provides a 1 Gbps (1000 Base-SX) full-duplex Ethernet LAN connection with throughput on a standard shortwave multimode optical cable that conforms to the IEEE 802.3z standard. The adapter supports distances of 260 m for 62.5 micron Multi Mode Fiber (MMF) and 550 m for 50.0 micron MMF. AIX Network Install Manager (NIM) boot capability is supported with this adapter.

Note: For optimum performance, the adapter should be placed in a 64-bit PCI-X card slot.

The IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700) incorporates an LC type connector on the card. This new, smaller form factor connector is being used by the industry for the next generation of fiber optic networks. If connecting into an older, existing SC type connector network, an LC-SC 62.5 Micron Fiber Converter Cable (FC 2459) or LC-SC 50 Micron Fiber Converter Cable (FC 2456) is required.

Configuration notes

Half-duplex (HDX) mode is not supported.

Hardware requirements

- ▶ One available PCI or PCI-X card slot

Software requirements

- ▶ AIX 5L Version 5.1 with the December 2002 or later Update CD or AIX 5L Version 5.2 with the December 2002 or later Update CD

9.7.9 IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701)

The IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter is a full-duplex Gigabit Ethernet adapter designed with highly integrated components. This adapter can be configured to run at 10, 100, or 1000 Mbps data rates. The adapter interfaces to the system using the PCI-X bus and connects to the network using a 4-pair CAT-5 Unshielded Twisted Pair (UTP) cable for distances of up to 100 m. AIX Network Install Manager (NIM) boot capability is supported with this adapter. The

adapter conforms to the IEEE 802.3ab 1000Base-T standard. The adapter also supports jumbo frames when running at the 1000 Mbps speed.

Note: For optimum performance, the adapter should be placed in a 64-bit PCI-X card slot.

Configuration notes

The 1000 Mbps speed is not supported in half-duplex (HDX) mode.

Hardware requirements

- ▶ One available PCI or PCI-X card slot

Software requirements

- ▶ AIX 5L Version 5.1 with the December 2002 or later Update CD or AIX 5L Version 5.2 with the December 2002 or later Update CD

9.7.10 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5706)

The IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter is a full-duplex, dual ported, Gigabit Ethernet adapter designed with highly integrated components. This adapter can be configured to run each port at 10, 100, or 1000 Mbps data rates. The adapter interfaces to the system using a PCI or PCI-X bus and connects to a network using a 4-pair CAT-5 Unshielded Twisted Pair (UTP) cable for distances of up to 100 m. AIX Network Install Manager (NIM) boot capability is supported with this adapter. The adapter conforms to the IEEE 802.3ab 1000Base-T standard. The adapter also supports jumbo frames when running at the 1000 Mbps speed. A function called 'Large Send', sometimes known as TCP Segmentation, is also provided by this adapter. This function offloads the TCP segmentation operation from the AIX IP layer to the adapter for outgoing (transmit side) TCP segments. Another function known as checksum offload, which offloads the TCP/UDP Checksum Operation or workload from the CPU to the adapter, is also provided. The IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5706) should be considered where maximum port density is required per I/O card slot. For a suggested maximum number of adapters, taking performance into consideration, refer to the *RS/6000 and IBM @server pSeries Adapter Placement Reference for AIX*, SA38-0538. If card slots are not the limiting factor and maximum throughput is required, the single port IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter (FC 5701) is the preferred solution.

Note: For optimum performance, the adapter should be placed in a 64-bit PCI-X card slot whenever possible.

Configuration notes

- ▶ The 1000 Mbps speed is not supported in half-duplex (HDX) mode.

Hardware requirements

- ▶ One available PCI or PCI-X card slot

Software requirements

- ▶ AIX 5L Version 5.1 or later or AIX 5L Version 5.2 or later

Note: For optimum 2-Port Gigabit Ethernet PCI-X Adapter performance, systems manufactured prior to December 6, 2002 may require a system firmware update. Check the following URL after December 6, 2002 to review and download latest firmware if needed:

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

9.7.11 IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707)

The IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter provides two 1 Gbps (1000 Base-SX) full-duplex Ethernet LAN connections with throughput on a standard shortwave multimode optical cable that conforms to the IEEE 802.3z standard. The adapter supports distances of 260 m for 62.5 micron Multi Mode Fiber (MMF) and 550 m for 50.0 micron MMF. AIX Network Install Manager (NIM) boot capability is supported with this adapter. A function called 'Large Send', sometimes known as TCP Segmentation, is also provided by this adapter. This function offloads the TCP segmentation operation from the AIX IP layer to the adapter for outgoing (transmit side) TCP segments. Another function known as checksum offload, which offloads the TCP/UDP Checksum Operation or workload from the CPU to the adapter, is also provided. The IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter (FC 5707) should be considered where maximum port density is required per I/O card slot. For a suggested maximum number of adapters, taking performance into consideration, refer to the *RS/6000 and IBM @server pSeries Adapter Placement Reference for AIX, SA38-0538*. If card slots are not the limiting factor and maximum throughput is required, the single port IBM Gigabit Ethernet-SX PCI-X Adapter (FC 5700) is the preferred solution.

Note: For optimum performance, the adapter should be placed in a 64-bit PCI-X card slot whenever possible.

Note: The 2-Port IBM Gigabit Ethernet-SX PCI-X Adapter incorporates an LC type connector on the card. This new, smaller form factor connector is being used by the industry for the next generation of fiber optic networks. If connecting into an older, existing SC type connector network, an LC-SC 62.5 Micron Fiber Converter Cable (FC 2459) or LC-SC 50 Micron Fiber Converter Cable (FC 2456) is required.

Configuration notes

- ▶ Half-duplex (HDX) mode is not supported.

Hardware requirements

- ▶ One available PCI or PCI-X card slot

Software requirements

- ▶ AIX 5L Version 5.1 or later or AIX 5L Version 5.2 or later

Note: For optimum 2-Port Gigabit Ethernet PCI-X Adapter performance, systems manufactured prior to December 6, 2002 may require a system firmware update. Check the following URL after December 6, 2002 to review and download latest firmware if needed:

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

9.7.12 2 Gigabit Fibre Channel PCI-X Adapter (FC 6239)

The 2 Gigabit Fibre Channel PCI-X Adapter is a 64-bit address/data, short form factor PCI-X adapter with LC type external fiber connectors that provides single or dual initiator capability over an optical fiber link or loop. With the use of appropriate optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage. The 2 Gigabit Fibre Channel PCI-X Adapter will auto-negotiate for the highest data rate (either 1 Gbps or 2 Gbps) of which the device or switch is capable. Distances of up to 500 meters running at 1 Gbps data rate and up to 300 meters running at 2 Gbps data rate are supported between the adapter and an attaching device or switch. When used with IBM supported Fibre Channel storage switches supporting long-wave optics, distances of up to 10 kilometers are capable running at either 1 Gbps or 2 Gbps data rates.

The 2 Gigabit Fibre Channel PCI-X Adapter can be used to attach devices either directly, or by means of Fibre Channel Switches. If attaching a device or switch with a SC type fiber connections), use of an LC-SC 50 Micron Fiber Converter

Cable (FC 2456) or a LC-SC 62.5 Micron Fiber Converter Cable (FC 2459) is required.

Refer to the following IBM storage subsystem Web page for additional supported server attachment information for IBM devices:

<http://www.storage.ibm.com/hardsoft/disk/products.htm>

Hardware requirements

- ▶ One empty PCI or PCI-X slot

Software requirements

- ▶ AIX 5L Version 5.1 or later or AIX 5L Version 5.2 or later

Note: Fibre Channel boot capability is supported on your system with this adapter, provided your system firmware is at the proper support level.

9.8 Cryptographic adapters

The following section describes the available cryptographic adapters.

9.8.1 IBM e-business Cryptographic Accelerator (FC 4960)

The IBM e-business Cryptographic Accelerator is a short form factor PCI Secure Socket Layer (SSL) hardware accelerator adapter. For Secure Web transactions, SSL operations is a key requirement. To do this, public-key cryptographic operations using SSL handshake protocol is employed. The IBM e-business Cryptographic Accelerator is a hardware cryptographic solution that off-loads this compute-intensive public-key cryptographic processing from the host.

The overall operation control, including command decoding, is implemented in hardware and requires no on-card microprocessor subsystem. As such, the adapter is a less expensive alternative to those who do not need the high security of the on-card secure programming environment, such as offered by the PCI Cryptographic Coprocessor (FC 4958), but do need the high cryptographic performance that hardware acceleration provides by offloading the host processor.

Note: The IBM e-business Cryptographic Accelerator is only supported by the industry-standard PKCS #11 application programming interface (API) Version 2.01, and applications which interface to the PKCS #11 Support Program.

Hardware requirements

One free PCI slot

Software requirements

AIX 5L Version 5.1 or later

9.8.2 PCI Cryptographic Coprocessor-FIPS-4 (FC 4963)

The PCI Cryptographic Coprocessor (FIPS-4) is a 2/3 length PCI adapter that combines hardware and software to provide a wide variety of security services. The PCI Cryptographic Coprocessor (FIPS-4) is a second generation adapter of the Cryptographic Coprocessor family that provides high performance secure hardware engines for secure Internet transactions, such as transmitting data, verifying electronic signatures, and bulk data encryption and decryption. In addition, the card is enclosed in a tamper proof enclosure to restrict access to on card resources, and is designed to FIPS 140-1 Level 4 standards.

Security functions supported by the adapter include:

- ▶ DES (Data Encryption Standard) (40-and 56-bit key) encryption and decryption, with pre-and post-padding. The coprocessor uses both ECB (electronic and codebook) and CBC (cipher block chain) modes of encryption.
- ▶ MAC (Message Authentication) generation and MAC verification services.
- ▶ Triple DES (three key) encryption and decryption of eight-byte units.
- ▶ Secure RSA key-pair generation.
- ▶ RSA signature generation and signature verification at 18 signatures/second at 2048 bits.
- ▶ Hardware random number generation.
- ▶ Secure data storage and retrieval.
- ▶ Other non-cryptographic security utilities can be carried out using the onboard processor.

As part of this adapter's physical security features, the following events will cause an adapter shutdown and secure data serialization:

- ▶ Shipping/Storage Temperature less than -15° C or greater than 95° C
- ▶ Dead battery (VBAT less than 2.4V)
- ▶ Supply voltage greater than 3.3V/12V max
- ▶ Mesh Sensor opens/shorts detection
- ▶ X-ray exposure

IBM offers software to enable your use of the Coprocessors. Two different approaches to cryptographic functions are offered for download from the following Web site:

<http://www.ibm.com/security/cryptocards>

- ▶ PKCS #11 Version 2.01, an implementation of the industry-standard API
- ▶ IBM Common Cryptographic Architecture (CCA), featuring support of special interest to the finance industry.

Under custom contract, IBM also offers toolkits that you can employ to develop extensions to the CCA offering and to develop your own application to exploit the secure computing environment and cryptographic hardware. For more information on custom contracts, refer to Web site at:

<http://www.ibm.com/security/cryptocards>

For additional information on the IBM PCI Cryptographic Coprocessor, refer to the following Web site:

<http://www.ibm.com/security/cryptocards>

Configuration notes

The IBM PCI Cryptographic Coprocessor Adapter is a field only installed device in order to meet restrictive shipping requirements.

Hardware requirements

One free PCI slot

Software requirements

AIX 5L Version 5.1 or later



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 pSeries workstations that your customer chooses. These accelerators support the X Window System and provide acceleration for key 2D functions and operations. In addition, most of IBM's family of 2D graphics accelerators 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, IBM's 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, we will discuss the graphics accelerators.

10.1 Available graphics accelerators

This section discusses the available graphics accelerators in detail.

10.1.1 POWER GXT135P graphics accelerator (FC 2848)

The POWER GXT135P Graphics PCI Adapter (Figure 10-1) is a high-performance, 2-D PCI graphics adapter that accelerates and enhances your system unit video. The POWER GXT135P Graphics PCI Adapter has no hardware switches to set. Mode selection is made through the software. Connection to the video monitor is made through a high-density 15-pin D-shell connector or a 28-pin DVI connector.




Figure 10-1 POWER GXT135P

The features of POWER GXT135P can be found in Table 10-1.

Table 10-1 Features of POWER GXT135P

Feature	Description
16-MB SDRAM	-
32-bit Universal PCI bus interface	-
8-bit or 24-bit color support	-
One or two analog monitors (identical image) supported at resolutions	<ul style="list-style-type: none">▶ 640 x 480 at 60 Hz▶ 1024 x 768 at 60 - 85 Hz▶ 1280 x 1024 at 60 - 85 Hz

Feature	Description
One analog monitor supported at resolutions	<ul style="list-style-type: none"> ▶ 1600 x 1200 at 75 - 85 Hz ▶ 2048 x 1536 at 60 - 75 Hz
One digital monitor (FC 2849 only) supported at resolutions	<ul style="list-style-type: none"> ▶ 640 x 480 at 60 Hz ▶ 1024 x 768 at 60 Hz ▶ 280 x 1024 at 60 Hz ▶ 600 x 1200 at 30 Hz

Hardware requirement

POWER GXT135P supported systems are described in Table 10-2.

Table 10-2 POWER GXT135P supported systems

System	Model
pSeries	<ul style="list-style-type: none"> ▶ pSeries 610-6C1/6E1 ▶ pSeries 615-6C3/6E3 ▶ pSeries 620 ▶ pSeries 640 ▶ pSeries 630-6C4/6E4 ▶ pSeries 650 ▶ pSeries 655 ▶ pSeries 660-6M1 ▶ pSeries 670 ▶ pSeries 690
RS/6000	<ul style="list-style-type: none"> ▶ 7043-150 ▶ 7044-170 ▶ 7044-270

Software requirements

The POWER GXT135P Graphics PCI Adapter is supported on AIX Version 4.3 with service package 4330-10 or later, AIX 5L Version 5.1 with service package 5110-01 or later, or AIX 5L Version 5.2. Support for digital monitors (FC 2849 only) requires AIX 5L Version 5.1 with service package 5100-04 or AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package.

10.1.2 POWER GXT4000P graphics accelerator (FC 2826)

The POWER GXT4000P 3D Graphics Adapter (Figure 10-2) is a single-card graphics adapter that attaches to your system unit in a PCI bus graphics slot. This adapter provides 3D graphics acceleration. Connection to the video monitor is made through a Digital Video Interface (DVI) connector.




Figure 10-2 POWER GXT4000P

The GXT4000P:

- ▶ Is designed to meet the price/performance needs of the value-oriented market for entry 3D workstations
- ▶ Shows significant improvement over previous IBM graphics accelerators, while providing a very attractive price point

The IBM Power GXT4000P Graphics Accelerator is a 64-bit entry 3D PCI Graphics adapter. This adapter has the base features shown in Table 10-3.

Table 10-3 Features of POWER GXT4000P

Feature	Description
128 MB Unified Frame Buffer	<ul style="list-style-type: none">▶ 24-bit Double Buffered up to 1920 x 1200▶ 24-bit Z-Buffer▶ 4/8-bit Overlay double-buffered▶ 8-bit Double Buffered Alpha▶ 8-bit Stencil/Clip Planes▶ Eight Windows ID bits

Feature	Description
One Rectangular Scissor region	-
Five HW Rectangular Clippers	-
Texture Mapping	<ul style="list-style-type: none"> ▶ Up to 108 MB Texture Memory (1280 x 1024) ▶ Dual Texture ▶ 3D Texture ▶ Texture Color Table
Video Support	<ul style="list-style-type: none"> ▶ Bilinear Scaling ▶ Color Space Conversion ▶ Four Hardware Color Maps ▶ Gamma Corrected AA lines ▶ HW occlusion culling
API Support	<ul style="list-style-type: none"> ▶ X11 ▶ GraPHIGS ▶ OpenGL 1.2
Monitor Support:	<ul style="list-style-type: none"> ▶ Resolutions supported: <ul style="list-style-type: none"> – 1280 x 1024 at 85 Hz – 1600 x 1200 at 85 Hz – 1920 x 1200 at 76 Hz ▶ DDC2B Support ▶ ISO 9241-Compliant

Hardware requirements

POWER GXT4000P supported systems are describe in Table 10-4.

Table 10-4 POWER GXT4000P supported systems

System	Model
RS/6000	<ul style="list-style-type: none"> ▶ 7043-150 ▶ 7044-170 ▶ 7044-270 ▶ 9112-265 ▶ 9114-275

Software requirements

The POWER GXT4000P Graphics PCI Adapter is supported on AIX Version 4.3.3 with AIX Version 4.3.3.0-05 Recommended Maintenance package or later. If support is required on other versions or releases of AIX, ensure that this adapter is supported on that release of AIX prior to installation.

10.1.3 POWER GXT4500P graphics accelerator (FC 2842)

The GXT4500P is a continuation of the successful family of graphics accelerators in support of pSeries and RS/6000 workstations. This new product is based on a high-performance raster engine, which is the result of a design implemented in state-of-the-art IBM chip technology.

The GXT4500P is designed to meet the price/performance needs of the value-oriented market for entry 3D workstations.

Figure 10-3 shows the POWER GXT4500P.




Figure 10-3 POWER GXT4500P

This new graphics accelerator, when configured in the pSeries POWER3-II-based workstations, offer an excellent combination of function and performance for complex and demanding graphics applications in such areas as:

- ▶ Mechanical Computer Aided Design (MCAD) and Engineering (MCAE) for automotive and aerospace
- ▶ Petroleum exploration and production
- ▶ Scientific visualization
- ▶ Other technical design and visualization

The GXT4500P accelerator delivers native support in hardware for the OpenGL and IBM graPHIGS application programming interfaces (APIs). This results in unmatched flexibility of support for application software, regardless of the graphics API chosen for implementation.

The IBM Power GXT4500P Graphics Accelerator is a 64-bit entry 3D PCI Graphics adapter. This adapter has the base features shown in Table 10-5.

Table 10-5 Features of POWER GXT4500P

Feature	Description
128 MB Unified Frame Buffer	<ul style="list-style-type: none"> ▶ 24-bit Double Buffered up to 2048 x 1536 on analog or digital displays ▶ 24-bit Double Buffered Stereo up to 1280 x 1024 ▶ 24-bit Z-Buffer ▶ 4/8-bit Overlay ▶ 8-bit Double Buffered Alpha ▶ 8-bit Stencil ▶ Eight Windows ID bits ▶ Four Clipping Planes
Scissor registers	-
Texture Mapping	<ul style="list-style-type: none"> ▶ Up to 110 MB Texture Memory (1280 x 1024) ▶ Dual Texture ▶ 3D Texture
Four Hardware Color Maps	-
API Support	<ul style="list-style-type: none"> ▶ X11 ▶ GraPHIGS ▶ OpenGL 1.2
Monitor Support	<ul style="list-style-type: none"> ▶ Resolutions supported: <ul style="list-style-type: none"> – 1024 x 768 at 85 Hz (analog) or 60 Hz (digital) – 1280 x 1024 at 85 Hz (analog) or 60 Hz (digital) – 1600 x 1200 at 85 Hz (analog) or 60 Hz (digital) – 1920 x 1200 at 76 Hz – 2048 x 1536 at 60 Hz (analog) or 30 Hz (digital) ▶ DDC2B Support ▶ ISO 9241-3- and ISO13406-2-Compliant

Hardware requirements

POWER GXT4500P supported systems are shown in Table 10-6.

Table 10-6 POWER GXT4500P supported systems

System	Model
pSeries	pSeries 630-6C4/6E4
RS/6000	<ul style="list-style-type: none">▶ 7043-150▶ 7044-170▶ 7044-270▶ 9112-265▶ 9114-275

Software requirements

The POWER GXT4500 is supported on AIX Version 4.3 with maintenance package 4330-10 or later. If support is required on an AIX release other than AIX Version 4.3 or later, ensure that this adapter is supported on that release of AIX prior to installation.

10.1.4 POWER GXT6000P graphics accelerator (FC 2827)

The POWER GXT6000P 3D Graphics Adapter is a single-card graphics adapter that attaches to your system unit in a PCI bus graphics slot. This adapter provides 3D graphics acceleration. Connection to the video monitor is made through a Digital Video Interface (DVI) connector.

The GXT6000P, with its hardware geometry acceleration:

- ▶ Supports the higher performance requirements of the more demanding user of advanced 3D workstations
- ▶ Provides the highest level of graphics performance available for the RS/6000 workstations
- ▶ Delivers significant advances in performance as measured by the Graphics Performance Characterization (GPC) benchmarks and, more importantly, within MCAD and MCAE applications

In some benchmarks, for certain environments, the GXT6000P more than doubles the speed of previous IBM graphics products.

In the GXT6000, the geometry engine performs very high precision transformations on geometric models and provides clipping and multisource lighting.

The IBM Power GXT6000P Graphics Accelerator is a 64-bit mid-range 3D PCI Graphics adapter. This adapter has the base features shown in Table 10-7.

Table 10-7 Features of POWER GXT6000P

Feature	Description
128 MB Unified Frame Buffer	<ul style="list-style-type: none"> ▶ 24-bit Double Buffered up to 1920 x 1200 ▶ 24-bit Quad Buffered Stereo up to 1280 x 1024 ▶ 24-bit Z-Buffer ▶ 4/8-bit Overlay ▶ 8-bit Double Buffered Alpha ▶ 8-bit Stencil/Clip Planes ▶ Eight Windows ID bits
One Rectangular Scissor Region	-
Five HW Rectangular Clippers	-
Texture Mapping	<ul style="list-style-type: none"> ▶ Up to 108 MB Texture Memory (1280 x 1024) ▶ Dual Texture ▶ 3D Texture ▶ Texture Color Table
Video Support	<ul style="list-style-type: none"> ▶ Bilinear Scaling ▶ Color Space Conversion
Four Hardware Color Maps	-
Gamma Corrected AA lines	-
Full OpenGL/graPHIGS Geometry Accelerator	-
HW occlusion culling	-
API Support	<ul style="list-style-type: none"> ▶ X11 ▶ GraPHIGS ▶ OpenGL 1.2
Monitor Support	<ul style="list-style-type: none"> ▶ Resolutions supported: <ul style="list-style-type: none"> – 1280 x 1024 at 85 Hz – 1600 x 1200 at 85 Hz – 1920 x 1200 at 76 Hz ▶ DDC2B Support ▶ ISO 9241-Compliant

Hardware requirement

POWER GXT6000P supported systems are shown in Table 10-8.

Table 10-8 POWER GXT6000P supported systems

System	Model
RS/6000	<ul style="list-style-type: none">▶ 7044-170▶ 7044-270▶ 9112-265▶ 9114-275

Software requirements

The POWER GXT6000P Graphics PCI Adapter is supported on AIX Version 4.3.3 with AIX Version 4.3.3.0-05 Recommended Maintenance package or later. If support is required on an other version or release of AIX, ensure that this adapter is supported on that release of AIX prior to installation.

10.1.5 POWER GXT6500P graphics accelerator (FC 2843)

The GXT6500P is a continuation of the successful family of graphics accelerators in support of RS/6000 workstations. These two new products are based on a high-performance raster engine, which is the result of a design implemented in state-of-the-art IBM chip technology.

The GXT6500P, with its hardware geometry accelerator chip implemented in copper technology, supports the higher performance requirements of the more demanding user of advanced 3D workstations.

The GXT6500P accelerator delivers native support in hardware for the OpenGL and IBM graPHIGS application programming interfaces (APIs). This results in unmatched flexibility of support for application software, regardless of the graphics API chosen for implementation.

Figure 10-4 on page 426 shows POWER GXT4500P.




Figure 10-4 POWER GXT6500P

In the GXT6500P, the geometry engine performs very-high, precision transformations on geometric models and provides clipping and multisource lighting.

The IBM Power GXT6500P Graphics Accelerator is a 64-bit, mid-range, 3D PCI Graphics adapter. This adapter has the base features shown in Table 10-9.

Table 10-9 Features of POWER GXT6500P

Feature	Description
128 MB Unified Frame Buffer	<ul style="list-style-type: none"> ▶ 24-bit Double Buffered up to 2048 x 1536 on analog or digital displays ▶ 24-bit Double Buffered Stereo up to 1280 x 1024 ▶ 24-bit Z-Buffer ▶ 4/8-bit Overlay ▶ 8-bit Double Buffered Alpha ▶ 8-bit Stencil ▶ Eight Windows ID bits ▶ Four Clipping Planes
Scissor registers	-
Texture Mapping	<ul style="list-style-type: none"> ▶ Up to 110 MB Texture Memory (1280 x 1024) ▶ Dual Texture ▶ 3D Texture
Four Hardware Color Maps	-
Full OpenGL/graPHIGS Geometry Accelerator	-

Feature	Description
API Support	<ul style="list-style-type: none"> ▶ X11 ▶ GraPHIGS ▶ OpenGL 1.2.1
Monitor Support	<ul style="list-style-type: none"> ▶ Resolutions supported, 8/24 bits: <ul style="list-style-type: none"> – 1024 x 768 at 85 Hz (analog) or 60 Hz (digital) – 1280 x 1024 at 85 Hz (analog) or 60 Hz (digital) – 1600 x 1200 at 85 Hz (analog) or 60 Hz (digital) – 1920 x 1200 at 76 Hz – 2048 x 1536 at 60 Hz (analog) or 30 Hz (digital) ▶ DDC2B Support ▶ ISO 9241-3- and ISO13406-2-Compliant

Hardware requirement

POWER GXT6500P supported systems are shown in Table 10-10.

Table 10-10 POWER GXT6500P supported systems

System	Model
pSeries	pSeries 630-6C4/6E4
RS/6000	<ul style="list-style-type: none"> ▶ 7044-170 ▶ 7044-270

Software requirements

The POWER GXT6500P is supported on AIX Version 4.3 with maintenance package 4330-10 or later. If support is required on an AIX release other than AIX Version 4.3 or later, ensure that this adapter is supported on that release of AIX prior to installation.

10.1.6 Display and cable matrix

Table 10-11 on page 428 shows the cable feature to graphics accelerator reference. Check your specific system model to determine which graphics accelerators are supported on which model.

Table 10-11 Cable feature to graphics accelerator

Display	GXT130P/ GXT110P (5D/DDC) (1)	GXT135P	GXT300P (15D/DDC)	GXT2000P/ GXT3000P (15D/DDC)	GXT4000P/ GXT6000P (DVI/DDC)	GXT4500P/ GXT6500 (DVI/DDC)
P70	4238	4238	4238	4238	4238(6)	4238(6)
P72(2)	Display	Display	Display	Display	Display(6)	Display(6)
P76(2)	Display	Display	Display	Display	Display(6)	Display(6)
P77(2)	Display(3)	Display	Display	Display	Display(6)	Display(6)
P92(4)	Display	Display	Display	Display	Display(6)	Display(6)
P200	4238	4238	4238	4238	4238(6)	4238(6)
P201	4237	4237	4237	4237	4237(6)	4237(6)
P202(4)	Display	Display	Display	Display	Display(6)	Display(6)
P260(4)	Display	Display	Display	Display	Display	Display
P275(4)	Display(3)	Display	Display	Display	Display	Display
6091-19I	NS(5)	NS	NS	4239(7)	NS	NS
POWERDP17	4217	NS	4217	4239(7)	NS	NS
POWERDP20	4217	NS	4217	4239(7)	NS	NS
9516-B03	4217	4217	4217	4217	4217(6)	4217(6)
9511-Axx T54A/T54H/ T540	Display	Display	NS	NS	NS	NS
9519-Axx T85A	Display	Display	NS	Display	NS	Display
9493-Axx T56A	Display	Display	NS	NS	NS	NS
9494-xxx T860	NS	Display	NS	NS	NS	Display
9497-Axx T86A	Display(3)	Display	Display	Display	NS	Display(6)
6656-Hxx T560	NS	Display	NS	NS	NS	NS

Display	GXT130P/ GXT110P (5D/DDC) (1)	GXT135P	GXT300P (15D/DDC)	GXT2000P/ GXT3000P (15D/DDC)	GXT4000P/ GXT6000P (DVI/DDC)	GXT4500P/ GXT6500 (DVI/DDC)
6658-Hxx T84H	Display(3)	Display	NS	Display	NS	Display
6659-Hxx T210	NS	Display	NS	NS	NS	Display

Notes:

1. GXT130P only.
2. Captured 15-pin, D-shell cable.
3. Display: The appropriate cable is included with the display.
4. Dual input on monitor.
5. Not Supported: This display/adaptor combination is not supported.
6. Must use DVI to 15-pin D-shell converter supplied with the GXT4000P, GXT4500P, GXT6000P, or GXT6500P.
7. Feature number 4239 (3 BNC Cable) is recommended on GXT2000P and GXT3000P. Feature number 4217 is a 5 BNC cable (external in H and V sync).

10.2 Withdrawn graphics accelerators

Table 10-12 provides a list of the graphics accelerators that are withdrawn. The withdrawn dates are for the US.

Table 10-12 *Withdrawn graphics accelerator*

Description	Feature code	Withdrawn date
POWER Gt3™	2777	April 1993
High-Performance 8-Bit 3D Color Graphics Processor	2780	May 1993
High-Performance 24-Bit 3D Color Graphics Processor	2781	May 1993
POWER Gt4x™ 8-Bit	2790	December 1993
POWER Gt4x 24-Bit	2791	December 1993
POWER Gt4™ 8-Bit	2795	December 1993
POWER Gt4 24-Bit	2796	December 1993

Description	Feature code	Withdrawn date
Grayscale Graphics Display Adapter	2760	June 1994
POWER GTO™ Accelerator	4350	November 1994
Color Graphics Display Adapter	2770	January 1995
POWER Gt1x	4207	September 1995
POWER Gt1™	4208	September 1995
POWER Gt3i™	2768	September 1995
IBM E15-type Graphic	2731	January 1996
POWER Gt4e™	2776	October 1996
S15 Graphics Accelerator	2657	January 1997
POWER GXT155L	2665	July 1997
POWER GXT150L™	2660	September 1997
POWER GXT500P	2854	January 1998
POWER GXT1000™	7252	March 1998
POWER GXT110P	2839	August 1998
POWER GXT1000	7250-001/002	July 1998
POWER GXT800M	2850	March 1999
MVP	2837	March 1999
POWER GXT255P	2852	September 1999
POWER GXT550P	2845	September 1999
POWER GXT550P	2855	September 1999
POWER GXT800P	2853	September 1999
POWER GXT800P	2859	September 1999
GXT120P graphics adapter	2838	September 2000
GXT250P graphics adapter	2851	January 2001
GXT2000P graphics adapter	2823	December 2001
GXT3000P graphics adapter	2825	December 2001
GXT300P graphics adapter	2841	December 2001

Description	Feature code	Withdrawn date
GXT4000P graphics adapter	2826	December 2002
GXT6000P graphics adapter	2827	December 2002
GXT130P graphics adapter	2830	April 2003
GXT135P graphics adapter	2848	December 2003



Internal storage devices

This chapter discusses the available internal storage devices, including:

- ▶ DVD and CD-ROM
- ▶ Internal tape drives
- ▶ SCSI disk drives

11.1 DVD and CD-ROM

This section covers the available DVD and CD-ROM devices.

4.7 GB SCSI-2 DVD-RAM Drive black bezel (FC 2623)

The SCSI-2 DVD-RAM drive is a multifunction storage device capable of reading and writing 4.7 GB DVD-RAM discs and reading CD-ROM discs. This device also provides read/write compatibility with conventional 2.6 GB DVD-RAM discs. System boot and install functions are supported with CD-ROM and DVD-RAM media.

Characteristics:

- ▶ Media Data Transfer Rate: 3600 KB/s CD-ROM (24X) max at outer diameter, 1350 KB/s write (1X DVD) and 2705 KB/s read (2X DVD)
- ▶ Interface: SCSI-2 supporting synchronous and asynchronous data transfer
- ▶ Average Random Access Time: 100 ms CD-ROM, 155 ms DVD-RAM
- ▶ High-speed synchronous burst rate of 10 MBps
- ▶ Loading tray accommodates both 8 cm discs (in horizontal orientation only), 12 cm discs (CD-ROM), and 4.7 GB DVD cartridges
- ▶ Operates in either vertical or horizontal position
- ▶ Reads multi-session discs
- ▶ Reads CD-recordable discs
- ▶ Reads CD-RW discs
- ▶ Reads and writes 2.6 GB, 4.7 GB, and 9.4 GB (double sided) DVD-RAM media
- ▶ Supports all major CD-ROM formats, including Mode 1, Mode 2, XA, CDDA, and audio
- ▶ Contains headphone output and line output for audio

Limitations:

- ▶ DVD video is not supported.
- ▶ Attributes provided: 4.7 GB 2X DVD (2705 KB/s) DVD-RAM.
- ▶ Attributes required: One half high media bay and one SCSI address.

32X (Max) SCSI-2 CD-ROM Drive (FC 2624)

The 32X (Max) SCSI-2 Internal CD-ROM Drive is a tray loading CD-ROM drive providing up to 4800 KB/s maximum media data transfer rate. It is a 5.25-in. 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.

Characteristics:

- ▶ Media Data Transfer Rate: 4800 KB/s (Max)
- ▶ Interface: SCSI-2 8-bit Single Ended (SE)
- ▶ Average Random Access Time: 90 ms (typical)
- ▶ Buffer Memory: 512 KB
- ▶ Media capacity: Greater than 600 MB
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 5.25-in. half high form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats
- ▶ Loading tray supports 12 cm and 8 cm disks
- ▶ Attributes provided: 600 MB 32X max. (up to 4800 KB/s) CD-ROM
- ▶ Attributes required: One half high media bay and one SCSI-2 internal SE 8-bit address

48X (Max) SCSI-2 Internal CD-ROM Drive (FC 2628)

The 48X (Max) SCSI-2 Internal Auto-docking CD-ROM Drive is a tray loading CD-ROM drive providing up to 7200 KB/s maximum media data transfer rate. It is a 5.25-in. 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.

Characteristics:

- ▶ Media Data Transfer Rate: 7200 KB/s (Max)
- ▶ Interface: SCSI-2 8-bit Single Ended (SE)
- ▶ Average Random Access Time: 90 ms (typical)
- ▶ Buffer Memory: 512 KB
- ▶ Media capacity: Greater than 600 MB
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 5.25-in. half high form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats

- ▶ Loading tray supports 12 cm and 8 cm disks
- ▶ Attributes provided: 48X (Max) SCSI-2 Internal Auto-docking CD-ROM Drive
- ▶ Attributes required: One half high Auto-docking media bay and one SCSI-2 internal SE 8-bit address

4.7 GB SCSI-2 Auto-docking DVD-RAM Drive (FC 2629)

The SCSI-2 DVD-RAM drive is a multifunction storage device capable of reading and writing 4.7 GB DVD-RAM discs and reading CD-ROM discs. This device also provides read/write compatibility with conventional 2.6 GB DVD-RAM discs. System boot and install functions are supported with CD-ROM and DVD-RAM media.

Characteristics:

- ▶ Media Data Transfer Rate: 3600 KB/s CD-ROM (24X) max at outer diameter, 1350 KB/s write (1X DVD) and 2705 KB/s read (2X DVD)
- ▶ Interface: SCSI-2 supporting synchronous and asynchronous data transfer
- ▶ Average Random Access Time: 100 ms CD-ROM, 155 ms DVD-RAM
- ▶ High-speed synchronous burst rate of 10 MBps
- ▶ Loading tray accommodates both 8 cm discs (in horizontal orientation only), 12 cm discs (CD-ROM) and 4.7 GB DVD cartridges
- ▶ Operates in either vertical or horizontal position
- ▶ Reads multi-session discs
- ▶ Reads CD-recordable discs
- ▶ Reads CD-RW discs
- ▶ Reads and writes 2.6 GB, 4.7 GB, and 9.4 GB (double sided) DVD-RAM media
- ▶ Supports all major CD-ROM formats, including Mode 1, Mode 2, XA, CDDA, and audio
- ▶ Contains headphone output and line output for audio

Limitations:

- ▶ DVD video is not supported.
- ▶ Attributes provided: 4.7 GB SCSI-2 Auto-docking DVD-RAM Drive.
- ▶ Attributes required: One half high Auto-docking media bay and one SCSI address.

IDE CD-ROM Drive black bezel (FC 2633)

The 48X (Max) IDE Internal CD-ROM Drive is a tray loading CD-ROM drive providing up to 7200 KB/s maximum media data transfer rate. It is a 5.25-in. half high form factor, multi-session capable, CD-ROM drive that provides state of the art performance and supports existing 650 MB CD-ROM discs.

Characteristics:

- ▶ Media Data Transfer Rate: 7200 KB/s (Max)
- ▶ Interface: IDE/ATAPI
- ▶ Average Random Access Time: 75 ms (typical)
- ▶ Buffer Memory: 128 KB
- ▶ Media capacity: 650 MB
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 5.25-in. half high form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats
- ▶ Loading tray supports 12 cm and 8 cm disks
- ▶ Attributes provided: 650 MB 48X (Max) (Up to 7200 KB/s) IDE CE-ROM
- ▶ Attributes required: One half high media bay

16X/48X (max) IDE DVD-ROM Drive (FC 2634)

The 16X/48X (max) IDE DVD-ROM Drive is an internal tray loading DVD-ROM drive providing up to 7200 KB/s (CD-ROM) and 22.16 MBps (DVD-ROM) data transfer rates. It is a 5.25-in. half high form factor multi-session capable, DVD-ROM drive that provides state of the art performance and supports existing 650 MB CD-ROM, 4.7 GB DVD-ROM, and 9.4 GB DVD-ROM (double-sided) discs. This drive also reads Type II (removable from cartridge) DVD-RAM discs at DVD-ROM speeds. System boot and install functions are supported with CD-ROM and DVD-RAM media.

Characteristics:

- ▶ Media Data Transfer Rate: CD-ROM: 7200 KB/s (max); DVD-ROM: 22.16 MBps (max)
- ▶ Interface: IDE/ATAPI
- ▶ Average Random Access Time: CD-ROM: 90 ms (typical); DVD-ROM: 135 ms (typical)
- ▶ Buffer Memory: 256 KB

- ▶ Media capacity: CD-ROM: 650 MB; DVD-ROM: 4.7 GB (single sided) or 9.4 GB (double-sided)
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 5.25-in. half high form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats
- ▶ Loading tray supports 12 cm and 8 cm disks

Limitations:

- ▶ DVD video is not supported.
- ▶ Attributes provided: 16X/48X (max) IDE DVD-ROM Drive.
- ▶ Attributes required: One half high media bay.

16X/48X (max) SCSI Auto-docking DVD-ROM Drive (FC 2635)

The 16X/48X (max) SCSI Auto-Docking DVD-ROM Drive is an internal tray loading DVD-ROM drive providing up to 7200 KB/s (CD-ROM) and 22.16 MBps (DVD-ROM) data transfer rates. It is a 5.25-in. half high form factor multi-session capable, DVD-ROM drive which provides state of the art performance and supports existing 650 MB CD-ROM, 4.7 GB DVD-ROM, and 9.4 GB DVD-ROM (double-sided) discs. This drive also reads Type II (removable from cartridge) DVD-RAM discs at DVD-ROM speeds. System boot and install functions are supported with CD-ROM and DVD-RAM media.

Characteristics:

- ▶ Media Data Transfer Rate: CD-ROM: 7200 KB/s (max); DVD-ROM: 22.16 MBps (max)
- ▶ Interface: SCSI
- ▶ Average Random Access Time: CD-ROM: 90ms (typical); DVD-ROM: 135 ms (typical)
- ▶ Buffer Memory: 256 KB
- ▶ Media capacity: CD-ROM: 650 MB; DVD-ROM: 4.7GB (single sided) and 9.4 GB (double-sided)
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 5.25-in. half high form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats
- ▶ Loading tray supports 12 cm and 8 cm disks

Limitations:

- ▶ DVD video is not supported.
- ▶ Attributes provided: 16X/48X(max) SCSI Auto-Docking DVD-ROM Drive.
- ▶ Attributes required: One half high Auto-docking media bay and one SCSI-2 address.

IDE Slimline DVD-ROM Drive (FC 2640)

The 8X/24X (max) Slimline IDE DVD-ROM Drive is an internal tray loading DVD-ROM drive providing up to 3600 KB/s (CD-ROM) and 10.3 MBps (DVD-ROM) data transfer rates. It is a 12.7 mm Slimline form factor multi-session capable, DVD-ROM drive that provides state of the art performance and supports existing 650 MB CD-ROM, 4.7 GB DVD-ROM, and 9.4 GB DVD-ROM (double-sided) discs. This drive also reads Type II (removable from cartridge) DVD-RAM discs at DVD-ROM speeds. System boot and install functions are supported with CD-ROM and DVD-RAM media.

Characteristics:

- ▶ Media Data Transfer Rate: CD-ROM: 3600 KB/s (max); DVD-ROM: 10.3 MBps (max)
- ▶ Interface: IDE/ATAPI
- ▶ Average Random Access Time: CD-ROM: 95 ms (typical); DVD-ROM: 150 ms (typical)
- ▶ Buffer Memory: 256 KB
- ▶ Media capacity: CD-ROM: 650 MB; DVD-ROM: 4.7 GB (single sided) and 9.4 GB (double-sided)
- ▶ Multisession capable (Reads CD-R and CD-RW media)
- ▶ 12.7 mm Slimline form factor
- ▶ Operates in either vertical or horizontal positions
- ▶ Interface supports standard and extended XA formats
- ▶ Loading tray supports 12 cm and 8 cm disks

Limitations:

- ▶ DVD video is not supported.
- ▶ DVD-ROM only reads CD-type formatted media with AIX 5L Version 5.1.
- ▶ Attributes provided: 8X/24X(max) IDE DVD-ROM Drive.
- ▶ Attributes required: One Slimline media bay.

11.2 Internal tape drives

The following are internal tape drives available on the pSeries.

IBM 80/160 GB Internal Tape Drive with VXA (FC 6120)

The IBM 80/160 GB Internal Tape Drive with VXA Technology is a 5.25-in., half high, Ultra2 LVD 16-bit tape drive, which provides a high capacity for save/restore and archive functions. This tape drive uses VXA tape data cartridges and is compression capable, providing a capacity of up to 160 GB - a significant increase in capacity over the previous internal tape drives.

Characteristics:

- ▶ Capacity: 80 GB native mode, 160 GB (typical) compression mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: uses VXA tape data cartridges
- ▶ Technology: Helical scan, rotating head
- ▶ Operation: Streaming
- ▶ Data Transfer Rate: 6 MBps native mode, 12 MBps (typical) compression
- ▶ Interface: SCSI-2 (LVD/SE) asynchronous/synchronous
- ▶ Compatibility: 80 GB mode (Read/Write), 160 GB compression (Read/Write)
- ▶ Attributes provided: One 80/160 GB internal tape drive
- ▶ Attributes required: One 1.6-in. (41 mm) half high media bay and one SCSI-2 internal 16-bit address

60/150 GB 16-bit 8 mm Internal Tape Drive (FC 6131)

The 60/150 GB 16-bit 8 mm Internal Auto-docking Tape Drive consists of a 5.25-in. half high, 16-bit tape drive. This drive provides a high capacity tape drive for save/restore and archiving functions. This tape drive uses IBM 8 mm data cartridges and is compression capable, providing a capacity of up to 150 GB - up to 375% increase over the previous 20/40 GB 8 mm internal tape drive.

Characteristics:

- ▶ Capacity: 60 GB Native Mode, 150 GB (typical) Compression Mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: IBM 8 mm Data Cartridge with Smart Clean Technology
- ▶ Technology: Helical Scan, Rotating Head
- ▶ Operation: Streaming

- ▶ Data Transfer Rate: 12 MBps Native Mode, 30 MBps (typical) Compression Mode
- ▶ Interface: SCSI-2 16-bit Low Voltage Differential (LVD)/Single-ended (SE) Asynchronous/Synchronous
- ▶ Compatibility: 20 GB Mode (Read only), 20 GB Compression (Read only), 60 GB Mode (RW), 150 GB Compression (RW)
- ▶ Attributes provided: 60/150 GB 16-bit 8 mm Internal Auto-docking Tape Drive
- ▶ Attributes required: One 1.6-in. (41 mm) half high auto-docking media bay and one SCSI-2 internal 16-bit address

60/150 GB 16-bit 8 mm Internal Tape Drive (FC 6134)

The 60/150GB 16-bit 8 mm Internal Tape Drive consists of a 5.25-in. half high, 16-bit tape drive. This drive provides a high capacity tape drive for save/restore and archiving functions. This tape drive uses IBM 8 mm data cartridges and is compression capable, providing a capacity of up to 150 GB - up to 375% increase over the previous 20/40 GB 8 mm internal tape drive.

Characteristics:

- ▶ Capacity: 60 GB Native Mode, 150 GB (typical) Compression Mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: IBM 8 mm Data Cartridge with Smart Clean Technology
- ▶ Technology: Helical Scan, Rotating Head
- ▶ Operation: Streaming
- ▶ Data Transfer Rate: 12 MBps Native Mode, 30 MBps (typical) Compression Mode
- ▶ Interface: SCSI-2 16-bit Low Voltage Differential (LVD)/Single-ended (SE) Asynchronous/Synchronous
- ▶ Compatibility: See the following Technical Support Web page (8 mm Tape Interchange Information section) for compatibility information:
<http://techsupport.services.ibm.com/server/mdownload/tapewhdr.html>
- ▶ Attributes provided: 60/150 GB 16-bit 8 mm Internal Tape Drive
- ▶ Attributes required: One 1.6-in. (41 mm) half high media bay and one SCSI-2 internal 16-bit address

20/40 GB 4 mm Internal Tape Drive (FC 6158)

The 20/40 GB 4 mm Internal Tape Drive is a 5.25-in., half high, single-ended 16-bit tape drive, which provides a high capacity for save/restore and achieve functions. This tape drive uses IBM 4 mm data cartridges and is compression

capable, providing a capacity of up to 40 GB - a significant increase in capacity over the previous 12/24 4 mm internal tape drives (when using DDS-4 media).

Characteristics:

- ▶ Capacity: 20 GB Native Mode, 40 GB (typical) Compression Mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: IBM 4 mm DDS-4 data cartridge
- ▶ Technology: Helical scan, rotating head
- ▶ Operation: Streaming
- ▶ Data Transfer Rate: 3 MBps Native Mode, 6 MBps (typical) Compression Mode
- ▶ Interface: SCSI-2 (single ended) asynchronous/synchronous
- ▶ Compatibility: 4 GB mode (Read/Write), 8 GB compression (Read/Write); 12 GB mode (Read/Write), 24 GB compression (Read/Write), 20 GB mode (Read/Write), 40 GB compression (Read/Write)

Note: MES orders are shipped with both a black and a white bezel.

- ▶ Attributes provided: 4 mm tape capability
- ▶ Attributes required: One 1.6-in. (41 mm) half high media bay and one SCSI-2 internal SE 16-bit address

IBM 80/160 GB Internal Tape Drive with VXA (FC 6169)

The IBM 80/160 GB Internal Auto-docking Tape Drive with VXA Technology is a 5.25-in., half high, Ultra2 LVD 16-bit tape drive, which provides a high capacity for save/restore and achieve functions. This tape drive uses VXAtape data cartridges and is compression capable, providing a capacity of up to 160 GB - a significant increase in capacity over previous internal tape drives.

Characteristics:

- ▶ Capacity: 80 GB Native Mode, 160 GB (typical) Compression Mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: Uses VXAtape data cartridge
- ▶ Technology: Helical scan, rotating head
- ▶ Operation: Streaming
- ▶ Data Transfer Rate: 6 MBps Native Mode, 12 MBps (typical) Compression Mode

- ▶ Interface: SCSI-2 (LVD/SE) asynchronous/synchronous
- ▶ Compatibility: 80 GB mode (Read/Write), 160 GB compression (Read/Write)
- ▶ Attributes provided: One 80/160 GB internal auto-docking tape drive
- ▶ Attributes required: One 1.6-in. (41 mm) half high auto-docking media bay and one SCSI-2 internal 16-bit address

20/40 GB 4 mm Internal Auto-docking Tape Drive (FC 6185)

The 20/40 GB 4 mm Internal Auto-docking Tape Drive is a 5.25-in., half high, single-ended 16-bit tape drive, which provides a high capacity for save/restore and archive functions. This tape drive uses IBM 4 mm data cartridges and is compression capable, providing a capacity of up to 40 GB - a significant increase in capacity over the previous 12/24 4 mm internal tape drives (when using DDS-4 media).

Characteristics:

- ▶ Capacity: 20 GB Native Mode, 40 GB (typical) Compression Mode
- ▶ Form Factor: 5.25-in. half high
- ▶ Media: IBM 4 mm DDS-4 data cartridge
- ▶ Technology: Helical scan, rotating head
- ▶ Operation: Streaming
- ▶ Data Transfer Rate: 3 MBps Native Mode, 6 MBps (typical) Compression Mode
- ▶ Interface: SCSI-2 (single ended) asynchronous/synchronous
- ▶ Compatibility: 4 GB mode (Read/Write), 8 GB compression (Read/Write); 12 GB mode (Read/Write), 24 GB compression (Read/Write), 20 GB mode Read/Write, 40 GB compression (Read/Write)

Note: MES orders are shipped with both a black and a white bezel.

- ▶ Attributes provided: 4 mm 20/40 GB internal auto-docking tape drive
- ▶ Attributes required: One 1.6-in. (41 mm) half high auto-docking media bay and one SCSI-2 internal SE 16-bit address

11.3 SCSI drives

SCSI (pronounced *scuzzy*) is the acronym for the Small Computer System Interface. Given the wide availability and generally low cost of devices that conform to the SCSI standard, SCSI has become one of the most common mechanisms in the open systems world to attach many kinds of IBM and non-IBM peripherals.

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, Ultra2SCSI, and Ultra3 SCSI nomenclature in official product names.

11.3.1 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.

11.3.2 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.

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.

Command tag queuing

Command tag queuing (CTQ) provides the ability to queue multiple commands to a SCSI-2 device. This improves throughput by using more of the SCSI bus bandwidth. For example, additional SCSI commands can be passed to a disk device while it is performing a seek from a previous read or write command.

11.3.3 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. 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.

11.3.4 Overview of SCSI-III standards

The breakdown of the SCSI-III standards is shown in Table 11-1.

Table 11-1 Breakdown of the SCSI-III standards

Standards	Abbreviation	Function
SCSI-III Interlocked Protocol	SIP	Describes the protocols used on the SIP 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

Standards	Abbreviation	Function
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
SCSI-III Stream Commands	SSC	Describes the commands used to transmit streams 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 operate at its full rated speed.

11.3.5 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 burst rate) of 20 MBps.
Fast-20	Fast-20 is a bus running at double the clock speed of Fast, or 20 MHz. Fast-20 typically refers to an 8-bit bus and can also be called Fast-20 Narrow. Fast-20 Wide is the 16-bit version, also known as Ultra SCSI.
Ultra SCSI	Ultra SCSI, as described above, is a subset of the SCSI-III specification. It is effectively a Fast SCSI bus running at 20 MHz. Ultra SCSI can produce a maximum transfer of 20 MBps over an 8-bit data path. Wide Ultra SCSI, the 16-bit version of Ultra SCSI, can transmit a maximum of 40 MBps.
Ultra2 SCSI	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).

Ultra320 SCSI Ultra320 SCSI is currently taking SCSI to new levels of performance. Named for its superior 320 MBps data transfer speed, the latest generation of SCSI technology incorporates the three management features of Ultra320 SCSI that specifically affect data transfer: CRC, domain validation, and double transition clocking. Ultra3 refers to the 160 MBps technology.

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. With SCSI differential, longer cable lengths, up to 25 meters, can be maintained.

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.

11.3.6 Available SCSI disk drives

Table 11-2 lists the available SCSI disk drives on pSeries.

Table 11-2 Available SCSI disk drives

Feature code	Disk drive
3153	18.2 GB Ultra3 SCSI 10K RPM Disk Drive/1 inch
3157	18.2 GB Ultra3 SCSI 10K RPM Disk Drive/1 inch
3158	36.4 GB Ultra3 SCSI 10K RPM Disk Drive/1 inch
3277	36.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly
3159	73.4 GB Ultra3 SCSI 10K RPM Disk Drive/1 inch
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly
3275	146.8 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly
3102	18.2 GB 10K RPM Ultra SCSI Disk Drive/1 inch
3263	18.2 GB 10K RPM Ultra3 SCSI Disk Drive Assembly
3119	36.4 GB 10K RPM Ultra3 SCSI Disk Drive/1 inch
3264	36.4 GB 10K RPM Ultra3 SCSI Disk Drive Assembly
3280	36.4 GB 15,000 RPM Ultra3 SCSI Disk Drive Assembly

Feature code	Disk drive
3118	18.2 GB 10K RPM Ultra SCSI Disk Drive/1 inch
3265	73.4 GB 10K RPM Ultra3 SCSI Disk Drive Assembly
3281	73.4 GB 15,000 RPM Ultra3 SCSI Disk Drive Assembly
3272	18.2 GB 10K RPM Ultra3 SCSI Disk Drive Assembly
3273	36.4 GB 10K RPM Ultra320 SCSI Disk Drive Assembly
3274	73.4 GB 10K RPM Ultra320 SCSI Disk Drive Assembly
3275	146.8 GB 10K RPM Ultra320 SCSI Disk Drive Assembly
3278	73.4 GB 15,000 RPM Ultra320 SCSI Disk Drive Assembly
3276	146.8 GB 10K RPM Ultra3 SCSI Disk Drive Assembly
3026	18.2 GB Ultra2 SCSI 10K RPM Disk Drive/1inch
3129	36.4 GB Ultra3 SCSI 10K RPM Disk Drive/1inch
3110	18.2 GB Enhanced Ultra SCSI Disk Drive/1 inch
3128	36.4 GB 10K RPM Ultra SCSI Enhanced Disk Drive/1 inch



External storage architectures and devices

The IBM Total Storage family of products is designed to deliver the quality and performance that customers expect from IBM, while leading the industry in open storage networking. Our mandate is to provide solutions that are based on open standards, and that are interoperable in today's heterogeneous environments. For additional information, see the following link:

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

The following storage types are discussed:

- ▶ External disk storage products
- ▶ Tape drive products
- ▶ Tape automation products
- ▶ SAN

12.1 IBM external disk storage

The following sections describe the following IBM Storage products:

- ▶ The IBM 7133 Serial Storage Architecture Disk System
- ▶ The IBM 2104 Expandable Storage Plus disk enclosure
- ▶ The IBM 2105 is the IBM TotalStorage® Enterprise Storage Server® (ESS)
- ▶ The IBM FASt Storage Servers

12.2 Disk systems summary

Table 12-1 summarizes the key features of the ESS 2105 Enterprise Storage Server, the IBM 7133 Serial Storage Architecture Disk System, and the The IBM 2104 Expandable Storage Plus disk enclosure.

Table 12-1 Features of the ESS 2105, 7133 D40/T40, and 2104 DU3/TU3T

External Disk	ESS 2105 800	7133 D40/T40	2104 DU3/TU3
Product	Enterprise Storage Server	Serial Disk System	EXP Plus
Platform Support	xSeries®, iSeries, AS/400®, pSeries, RS/6000, zSeries, S/390, OS/400®, AIX, Solaris, HP-UX, Dynix, DG/UX, OpenVMS, Tru64, Windows NT, Windows 2000, NetWare, Linux for S/390, z/OS®, zVM, OS/390, VM/ESA®, VSE/ESA™, MVS/ESA™, TPF, Linux commercial distributions ² , and SGI Origin IRIX	pSeries, select RS/6000 servers, xSeries, select Netfinity® servers, AIX, Windows NT, Windows 2000, Solaris, HP-UX, and many others	pSeries, select RS/6000 servers, and AIX
Host Connectivity	1 Gb and 2 Gb Fibre Channel/FICON™, ESCON, and SCSI	SSA, Fibre Channel ¹ , and SCSI ¹	Ultra3 SCSI
SAN Support	Direct, FC-AL, and Switched Fabric	FC-AL ¹ and Switched Fabric	N/A
Copy Services	FlashCopy® Version 1 and Version 2, PPRC Version 1 and Version 2	Remote Copy (up to 10km), 3-way copy, and Instant Copy ¹	N/A

External Disk	ESS 2105 800	7133 D40/T40	2104 DU3/TU3
Availability Features	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual controllers, concurrent microcode update, and dual-pathing driver	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual adapters, enclosure services, and dual-pathing driver	RAID with server adapters, redundant power/cooling, hot-swap drives, and enclosure services
Controller	SMB dual active; optional turbo feature	Single/multiple active adapters	Adapters in server
Cache (min, max)	8 GB, 64 GB	0, 32 MB write and 64 MB read	N/A
RAID Support	5, 10	0, 1, 5, 10 (using adapters), 1, and 10	0, 1, 5 (using server-based adapters)
Capacity (min, max)	582 GB, 55.9 TB (physical capacity)	72.8 GB, 14.0 TB	18.2 GB, 2.0 TB
Drive Interface	SSA	SSA	Ultra3 SCSI
Drive Support	18.2 GB, 36.4 GB, 72.8 GB, and 145.6 GB 10,000 rpm disk drives; 18.2 GB, 36.4 GB, and 72.8 GB 15,000 rpm disk drives	18.2 GB, 36.4 GB, 72.8 GB and 145.6 GB 10,000 rpm disk drives; 36.4 GB and 72.8 GB 15,000 rpm disk drives	18.2 GB, 36.4 GB, 73.4 GB, and 146.8 GB 10,000 rpm disk drives; 36.4 GB and 73.4 GB 15,000 rpm disk drives
Certifications	Microsoft RAID, Cluster and Data Center, GDPS®, HACMP, Novell NetWare, Linux	Microsoft RAID, Cluster, HACMP	HACMP

Table 12-2 summarizes the key features of the IBM FAStT Storage Servers.

Table 12-2 The key features of the IBM FAStT Storage Servers

External Disk	FAStT900	FAStT700	FAStT600	FAStT200
Machine /Model	1742/90 U	1742/1 RU	1722/60 U	3542/1 RU, 2 RU

Platform Support	pSeries, select RS/6000 servers, xSeries, select Netfinity servers, select Sun and HP UNIX servers and other Intel processor-based servers, Windows NT, Windows 2000, Netware, VMWare, Linux, AIX, Solaris, and HP-UX	pSeries, select RS/6000 servers, xSeries, select Netfinity servers, Windows NT, Windows 2000, NetWare, Linux, AIX, HP-UX, and Solaris	pSeries, xSeries, Windows 2000; optional support for AIX, Sun Solaris, HP-UX, Novell NetWare, and Linux	pSeries, select RS/6000 servers, xSeries, select Netfinity servers, AIX, Windows NT, Windows 2000, Dynix, HP-UX, Linux, NetWare, and Solaris
Host Connectivity	Fibre Channel	Fibre Channel	Fibre Channel	Fibre Channel
SAN Support	Direct, FC-AL, and Switched Fabric	Direct, FC-AL, and Switched Fabric	Direct, C-AL, and Switched Fabric	Direct, FC-AL, and Switched Fabric
Copy Services	Remote Copy and FlashCopy	Remote Copy and FlashCopy		FlashCopy
Availability Features	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual controllers, concurrent microcode update capability, and dual-pathing driver	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual controllers, concurrent microcode update capability, and dual-pathing driver	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual controllers, concurrent microcode update capability, and dual-pathing driver	Fault-tolerant, RAID, redundant power/cooling, hot-swap drives, dual controllers, concurrent microcode update capability, and dual-pathing driver
Controller	Dual active 2 Gb RAID controllers	Dual active 2 Gb RAID controllers	Dual active 2 Gb RAID controllers	Single/dual active
Cache (min, max)	2 GB and 2 GB	2 GB and 2 GB	256 MB and 256 MB	128 MB and 256 MB
RAID Support	0, 1, 3, 5, and 10	0, 1, 3, 5, and 10	0, 1, 3, 5, and 10	0, 1, 3, 5, and 10
Capacity (min, max)	32 GB and 32 TB	32 GB and 32 TB	72 GB and 6.1 TB	18.2 GB and 9.6 TB using EXP500 or EXP700
Drive Interface	2 Gb FC-AL	2 Gb FC-AL	2 Gb FC-AL	FC-AL

Drive Support	36.4 GB, 73.4 GB, and 146.8 GB 10,000 rpm disk drives; 18.2 GB, 36.4 GB, and 73.4 GB 15,000 rpm disk drives	36.4 GB, 73.4 GB, and 146.8 GB 10,000 rpm disk drives; 18.2 GB, 36.4 GB, and 73.4 GB 15,000 rpm disk drives	36.4 GB, 73.4 GB, and 146.8 GB 10,000 rpm disk drives; 18.2 GB, 36.4 GB, and 73.4 GB 15,000 rpm disk drives	36.4 GB, 73.4 GB, and 146.8 GB 10,000 rpm disk drives; 18.2 GB, 36.4 GB, and 73.4 GB 15,000 rpm disk drives
Certifications	Microsoft RAID, Cluster, NetWare Cluster, and VERITAS Clustering ⁴	Microsoft RAID, Cluster and Data Center, HACMP, NetWare Cluster, and VERITAS Clustering	Microsoft RAID, Cluster and Data Center, and VERITAS Clustering	Microsoft RAID, Cluster, HACMP, NetWare Cluster, and VERITAS Clustering

Notes:

1. Using IBM TotalStorage SAN Controller 160 (no cluster or HACMP support).
2. Red Hat, SuSE, or TurboLinux. Please verify specific product information for details.
3. Consult product information for details.
4. Also, verification will be completed for HP Service Guard.
5. The SAN Integration Server is an offering that has not yet been formally announced. IBM product plans and intent are subject to change and such products may never be made generally available.

12.2.1 IBM 7133 Serial Disk System Advanced Models T40 and D40

The following section describes the IBM 7133 Serial Disk system. Figure 12-1 on page 456 shows the IBM 7133 Serial Disk System Advanced Models T40 (deskside tower) and D40 (rack-mountable drawer).



Figure 12-1 IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer

Highlights

This section describe the highlights of the IBM 7133 Serial Disk system:

- ▶ Provides outstanding disk storage performance with advanced SSA bandwidth of 160 MBps
- ▶ New Ultra high speed 15,000 rpm hard disk drives, now available in capacities of 36.4 GB or 72.8 GB
- ▶ Provides high capacity up to 2.3 TB per tower or drawer and 14.0 TB per host adapter
- ▶ Enables disk sharing through simultaneous attachment of multiple UNIX, Windows NT, and Novell NetWare hosts to the IBM 7133 Serial Disk System
- ▶ Provides high availability with redundant data paths, redundant cooling units, and two power supplies
- ▶ Facilitates remote mirroring (up to 10 km connection distances) with the Advanced SSA Optical Extender
- ▶ Simplifies storage management when used with the IBM StorWatch™ Serial Storage Expert

The IBM TotalStorage 7133 Serial Disk System Advanced Models D40 and T40 provide highly available storage for UNIX, Windows NT, and Novell NetWare

servers. By implementing a powerful industry-standard serial technology, the 7133 Advanced Models D40 and T40 provide outstanding performance, availability, and attach ability.

The rack-mountable 7133 Advanced Model D40 drawer is designed for integration into a supported 19-in. rack. The 7133 Advanced Model T40 is a free-standing desk-side tower unit.

Both 7133 Advanced Models can be populated with 145.6 GB, 72.8 GB, 36.4 GB, and 18.2 GB 10,000 RPM and 72.8 GB and 36.4 GB 15,000 RPM disk drives. Drive capacities can be intermixed, providing the flexibility to build storage capacity from gigabytes to terabytes. The 7133 Advanced Models can be intermixed in the same loop with other models of the 7133 (Models 010, 020, 500 and 600) as well as with the 7131-4052.

To download the newest SSA filesets, adapter, and disk microcode, use the following link:

<http://www.storage.ibm.com/hardsoft/products/ssa/rs6k/index.html>

For the Rack-mountable D40/Desk-side T40 minimum to maximum capacity, see Table 12-3.

Table 12-3 Rack-mountable D40/Desk-side T40 minimum to maximum capacity

Disk	Capacity
18.2 GB disks (10,000 RPM)	72.8 GB to 291.2 GB
36.4 GB disks (10K RPM and 15K RPM)	145.6 GB to 582.4 GB
72.8 GB disks (10K RPM and 15K RPM)	291.2 GB to 1.16 TB
146.8 GB disks (10,000 RPM)	587.2 GB to 2.34 TB

Dimensions:

- ▶ D40: 171 mm x 444 mm x 737 mm
- ▶ T40: 610 mm x 210 mm x 861 mm

Weight:

- ▶ D40: 118.0 lb (53.5 kg)
- ▶ T40: 168.0 lb (76.2 kg)

Operating environment:

- ▶ Temperature: 50° to 104° F (10° to 40° C)
- ▶ Relative humidity: 8 to 80%

- ▶ Wet bulb maximum: 80.6 (23° C)
- ▶ Power requirement: 20 to 400 watts

12.2.2 Expandable Storage Plus 2104

The IBM TotalStorage 2104 Expandable Storage Plus (EXP Plus) provides flexible, scalable and low-cost disk storage for IBM RS/6000 and pSeries servers in a compact package. This new disk enclosure is ideal for enterprises, such as Internet or application service providers, that need high-performance external disk storage in a small footprint. Industry-standard Ultra3 SCSI interfaces combine with Ultra3 SCSI adapters to provide exceptional throughput performance of 160 MBps. In addition, RAID support is possible with IBM SCSI RAID adapters. With all of these capabilities, the EXP Plus provides an extremely cost-effective and leading-edge disk storage solution. Figure 12-2 shows the Rack Model DU3 and the Deskside tower TU3.



Figure 12-2 IBM Expandable Storage Plus SCSI

Highlights

This section describe the highlights of the IBM TotalStorage 2104 Expandable Storage Plus (EXP Plus).

- ▶ Incorporates high-performance Ultra3 SCSI disk storage with 160 MBps throughput

- ▶ Features up to fourteen 10,000 RPM (18.2 GB, 36.4 GB, 73.4 GB, and 146.8 GB) or 15,000 RPM (36.4 GB, 73.4 GB) disk drives
- ▶ Scales from up to 2 TB of capacity per drawer or tower to more than 28 TB per rack
- ▶ Provides single or split-bus configuration flexibility to one or two servers
- ▶ Facilitates high availability with hot-swappable disk drives, redundant power and cooling, and dual SCSI ports
- ▶ Enhances storage protection with RAID 0, 1, 1E, 5, and 5E options when used with IBM Ultra3 SCSI RAID adapters
- ▶ Facilitates high availability with hot-swappable disk drives, redundant power and cooling, and dual SCSI ports

High availability

To help ensure that information is accessible, the EXP Plus is designed for high availability. It features hot-swappable, auto-docking disk drives, as well as an optional redundant power supply and cooling unit, all designed to minimize downtime.

Availability is enhanced with the use of RAID options or UNIX software mirroring. SCSI RAID adapters provide RAID 0, 1, 1E, 5, and 5E functions and can be used to attach the EXP Plus to RS/6000 and pSeries servers for increased data protection. To further facilitate availability, SCSI enclosure services can monitor and report problems to the server.

The EXP Plus can be configured to meet a variety of configuration requirements. In single server configurations, up to fourteen drives can be installed on a single SCSI bus. Alternatively, the bus can be split into two separate SCSI buses with up to seven drives each. The second host port option provides Ultra3 connectivity to either the same server or a second server.

A single bus configuration with a second host port supports up to twelve disks and enables connectivity to two hosts for disk affinity in non-RAID environments. HACMP can then be used to provide server failover for high availability in non-concurrent mode using the EXP Plus with non-RAID SCSI host bus adapters.

Flexible and scalable storage

The EXP Plus helps meet the needs of many different types of environments. Two models are available to fit the need for stand-alone storage in one- or two-server environments, or for rack-mounted storage used by several servers.

The rack-mounted Model DU3 drawer can reside in a variety of 19-in. racks, including the IBM 7014-S00, 7014-T00, 7014-T42, 7015-R00, or the 2101-100. The largest rack (7014-T42) can hold up to fourteen EXP Plus DU3 drawers for a total capacity of more than 28 TB.

The stand-alone tower Model TU3 is suitable for one- or two-server environments, providing up to 2 TB of capacity in a small desk-side tower unit.

Both the rack-mounted DU3 and the desk-side tower TU3 can be populated with up to fourteen disk drives. Available options include 10,000 RPM (18.2 GB, 36.4 GB, 73.4 GB, and 146.8 GB) and 15,000 RPM (36.4 GB and 73.4 GB) disk drives. Drive capacities can be intermixed and drives can be added in increments as few as one or as many as thirteen. Housing multiple drawers in a rack provides the flexibility to build storage capacity from gigabytes to terabytes.

The EXP Plus can attach to RS/6000 servers by using a variety of SCSI adapters: SCSI-2 Fast and Wide, Ultra SCSI, Ultra2 SCSI, and Ultra3 SCSI. For the highest performance and availability, the enclosures can be combined with the IBM PCI 4-channel Ultra3 SCSI RAID Adapter (FC 2498), enabling exceptional Ultra3 SCSI 160 MBps throughput performance and multiple RAID options. Distances up to 20 meters are supported between disk enclosures and pSeries or RS/6000 servers using Ultra3 SCSI adapters.

Because the IBM NAS 300 is designed to support 24x7 operations, it incorporates a variety of high-reliability, high-availability features. These features include redundant hot-swap power supplies and hard disk drives, as well as dual RAID cards, Fibre Channel hubs, and engines. With this design, the IBM NAS 300 supports an organization's operating ability, even if a subsystem failure occurs.

Physical dimensions

Table 12-4 describe the physical dimensions of the IBM TotalStorage 2104 Expandable Storage Plus (EXP Plus).

Table 12-4 The physical dimensions of the IBM TotalStorage 2104 EXP Plus

DU3	TU3
444 mm/17.5 in. Wide	281 mm/11.0 in. Wide
552 mm/21.7 in. Depth	594 mm/23.5 in. Depth
125 mm/5.0 in. (3U) Height	529 mm/21.0 in. High
38.5 kg/85.0 lbs	54.5 kg/120 lbs

Supported adapters

The following describe the supported adapter for the IBM TotalStorage 2104 Expandable Storage Plus (EXP Plus):

- ▶ FC 2494 PCI 3-Channel Ultra2 SCSI RAID
- ▶ FC 2498 PCI 4-Channel Ultra3 SCSI RAID
- ▶ FC 6203 PCI Dual Channel Ultra3 SCSI
- ▶ FC 6205 PCI Dual Channel Ultra2 SCSI
- ▶ FC 6206 PCI Single-Ended Ultra SCSI
- ▶ FC 6208 PCI SCSI-2 Fast/Wide Single-Ended

For more information about supported servers and adapters, visit:

<http://www.storage.ibm.com/disk/expplus/supserver.htm>

12.2.3 Enterprise storage server 2105 Model 800

The IBM TotalStorage Enterprise Storage Server (ESS) Model 800 helps set new standards in performance, automation, and integration, as well as capabilities that support continuous availability to data for the on demand world. This storage system also supports many advanced functions, which can be critical for increasing data availability during planned outages and for protecting data from planned and unplanned outages. These advanced functions can provide important disaster recovery and backup protection.

Figure 12-3 on page 462 shows the Enterprise Storage Server 2105 Model 800.



Figure 12-3 The Enterprise Storage Server 2105 Model 800

Highlights

The next section describe the highlights of the Enterprise Storage Server 2105 Model 800:

- ▶ Supports storage sharing for IBM @server pSeries and UNIX, Microsoft Windows NT, Microsoft Windows 2000, Microsoft Windows Server 2003, Novell NetWare, Linux, and SGI IRIX platforms, IBM @server iSeries and AS/400 platforms, and IBM @server zSeries and S/390 platforms
- ▶ Designed to provide outstanding performance with dual cluster RISC SMP processors, large cache, and serial disk attachment
- ▶ Uses redundant hardware, mirrored write caches, and RAID-5 and RAID-10 protection for “disks” to support high availability for mission-critical business applications
- ▶ Supports fast data transfer rates with attached hosts using ESCON, FICON, Fibre Channel, 2 Gigabit Fibre Channel/FICON, or Ultra SCSI
- ▶ Can help increase administrative productivity by offering efficient, centralized operations management through a Web browser, CLI (Command Line Interface), or by using the ESS API
- ▶ Enables enterprises with multiple heterogeneous hosts to scale up to 55.9 TB physical disk capacity while maintaining excellent performance

Advanced copy functions for business continuance

The next section describes the advanced copy functions of the Enterprise Storage Server 2105 Model 800.

► FlashCopy Version 1

FlashCopy Version 1 provides an advanced, fast replication facility that can help significantly reduce application outages needed for backups and other copy applications. The FlashCopy Version 1 NOCOPY option, FlashCopy's efficient "copy on write NOCOPY option, can help reduce overhead while also allowing flexible reuse of disk capacity that would otherwise be dedicated to copy operations. With the NOCOPY option, rather than a physical byte-for-byte copy of the source volume, the only data copied to the target is that which is about to be changed or overlaid by the application.

► Flash Copy Version 2

The ESS now supports FlashCopy Version 2, which includes all the features of FlashCopy Version 1 as well as many enhancements designed to improve capacity management and utilization. Among these enhancements are:

Data Set FlashCopy This feature offers a new level of granularity for the z/OS environment, allowing more efficient use of the ESS capacity.

Data Set FlashCopy Allows the source and target copy to be different sizes and allows the copied data to reside at a different location in the volume.

Multiple Relationship FlashCopy

This function allows a volume to participate in multiple FlashCopy relationships (up to 12 simultaneous relationships) so that multiple copies of the same data can be made for testing, backup, and other applications. This feature offers increased flexibility and improved capacity management and utilization.

Incremental FlashCopy

Offers the ability to track and record changes that are made to the source and target volumes after the establishment of FlashCopy relationships. This allows the capability to refresh a LUN or volume to the source or target's point in time content using only the changed data. The refresh can occur in either direction. It also offers improved flexibility and faster FlashCopy completion times.

- ▶ Peer-to-Peer Remote Copy (PPRC) Version 1
PPRC Version 1 includes a synchronous remote data-mirroring technique designed to constantly maintain a current copy of the local application data on a remote site. PPRC-XD (Extended Distance) is included in PPRC Version 1, and provides an asynchronous remote copy function.
- ▶ Peer-to-Peer Remote Copy (PPRC) Version 2
PPRC Version 2 includes all of the functionality of PPRC Version 1 plus an asynchronous cascading function designed as a long-distance remote copy solution for open systems and IBM @server zSeries servers.
- ▶ Extended Remote Copy (XRC)
XRC is a copy function available for the z/OS and OS/390 operating systems. It maintains a copy of the data asynchronously at a remote location over unlimited distances, and is designed to provide the same high throughput and data protection regardless of the distance to the secondary site.
- ▶ High availability to safeguard data access
Support for 24 x7 operations is built into the ESS. The ability to implement RAID-5 and RAID-10 disk arrays helps provide data protection while remote copy technologies allow fast data backup and disaster recovery. The ESS features dual active processing clusters with failover switching, hot spares, hot-swappable disk drives, mirrored write cache and redundant power and cooling. The ESS also contains integrated proactive self-diagnostics to help prevent downtime by constantly monitoring system functions.

Physical characteristics

The following list describes the physical characteristics of the Enterprise Storage Server 2105 Model 800.

- ▶ Height:192 cm
- ▶ Width:139 cm
- ▶ Depth: 91 cm
- ▶ Weight:1204 kg

The following list describes the physical characteristics of the Enterprise Storage Server 2105 Model 800 with expansion frame.

- ▶ Height:192 cm
- ▶ Width: 294 cm
- ▶ Depth: 91 cm
- ▶ Weight: 2545 kg

Capacity upgrades

This will provide scalability to grow as your business grows/evolves. It provides investment protection by allowing you to add more capacity to your existing system. It is cost-effective, since you can increase your disk storage capacity without investing in a new system and can also mix larger capacity drives for more efficiency and cost effectiveness. It lowers your total cost of operations using better resource utilization, mixing larger capacity drives.

Capacities:

Feature 2122	18.2 GB disk 8-pack, 10K RP
Feature 2123	36.4 GB disk 8-pack, 10K RP
Feature 2124	72.8 GB disk 8-pack, 10K RP
Feature 2125	Disk eight-pack - 145.6 GB
Feature 2142	18.2 GB disk 8-pack, 15K RP
Feature 2143	36.4 GB disk 8-pack, 15K RP
Feature 2144	Disk eight-pack - 72.8 GB, 15K rpm

Cache upgrades:

Feature 4014	16 GB cache (8 GB additional on base)
Feature 4015	24 GB cache (16 GB additional on base)
Feature 4016	32 GB cache (24 GB additional on base)
Feature 4020	64 GB cache (56 GB additional on base)

12.2.4 IBM TotalStorage FASt900 Storage Server

The IBM 1742 TotalStorage FASt900 Storage Server Model 90U expands the family of Fibre Array Storage Technology (FASt) highly scalable offerings with improved performance capabilities. Built upon a fourth generation Fibre Channel (FC) RAID controller, the FASt900 brings high availability, advanced functionality, scalable capacity, and connectivity to a wide range of storage area network (SAN) applications in mission-critical enterprise networks. The FASt900 is designed to support applications running on IBM @server pSeries and xSeries servers as well as servers from Sun, Hewlett Packard, and other Intel-based providers.

Figure 12-4 on page 466 shows a FASt900 Model 90U.



Figure 12-4 The FASt900 Model 90U

Highlights

The following describe the highlights of the FASt900 storage server:

- ▶ Uses 2 Gb Fibre Channel connectivity to support high performance (772 MBps throughput from disk) for faster, more responsive access to data
- ▶ Facilitates storage consolidation for Storage Area Network (SAN), Network Attached Storage (NAS), and direct-attach storage environments
- ▶ Supports continuous availability through redundant components and mirrored, battery backed-up cache
- ▶ Supports two to 224 hard disk drives and provides more than 32 TB of easily scalable capacity
- ▶ Includes FASt Storage Manager to help manage the storage subsystem
- ▶ Provides flexibility for multiplatform storage environments by supporting a wide variety of servers, operating systems and cluster technologies
- ▶ Offers advanced replication functions such as FlashCopy and Remote Mirroring for high data availability and protection

Description

The FASt900 coupled with the EXP700 Expansion Unit offers you an industry-standard, highly available, high-performance storage solution that scales in capacity and performance within a wide range of SAN environments.

The IBM TotalStorage FASt family of storage products offers an ideal set of building blocks for your Fibre Channel FC-based storage area network (SAN), whether you are just starting out, or expanding your existing network. The FASt

product family is very scalable, from 36 GBs to over 32 TBs of capacity per FAStT Storage Server. It is easy to install and manage utilizing the FAStT Storage Manager, which is included with every FAStT Storage Server. The FAStT family of products offers various controller models at different performance and price points to best match your SAN needs. Multiple FAStT Storage Servers can be centrally managed from a common FAStT management console across your existing LAN infrastructure, and with the FAStT advanced copy service functions, such as Flashcopy and Remote Mirror, your data protection and disaster recovery needs can be addressed.

The key benefits of the FAStT family of storage products are:

- ▶ High availability: Dual processor and redundant hardware
- ▶ Scalability: From 36 GBs to over 32 TBs
- ▶ Performance: World class throughput to best match your applications needs
- ▶ Advanced functions for business continuance: FlashCopy and Remote Mirror

The FAStT900, the newest member of the FAStT family, offers higher performance capabilities with all of the high availability and scalability features you have come to expect with the FAStT. It is a 4U-rack mountable enclosure that contains dual RAID controllers and multiple SAN connectivity options used in conjunction with up to 16 of the FAStT EXP700 Expansion Units to provide a complete solution for your storage needs.

Key elements of the FAStT900 are:

- ▶ Dual, high-performance RAID controller cards, each with two 2 Gbps host ports
- ▶ Support for RAID 0, 1, 3, 5, and 10
- ▶ 2 GB of write-back cache, 1 GB per controller, with battery backup protection for up to three days
- ▶ Support for up to 224 disk drives with FAStT EXP700 Expansion Units, over 32 TB of capacity
- ▶ FAStT Storage Manager Version 8.3 with support for 16 storage partitions, option to expand to 64 storage partitions
- ▶ Redundant, hot-swappable power supplies and cooling fan assemblies
- ▶ Attachment to a wide range of industry-standard SAN switches and directors

The FAStT900 utilizes the latest 2 Gbps Fibre Channel technology for connections to host servers or switched SAN fabrics, and to disk drive enclosures. On the FAStT900 host interface, there are two host-side miniHUBs standard, each providing a 2 Gbps FC attachment to each of the RAID controllers. Optionally, up to two additional miniHUBs (Feature code 3507) can

be added, doubling the host attachment connectivity. Each host-side miniHUB can be directly attached to a switched fabric or to a host server's Fibre Channel host bus adapter using either shortwave Small Form-factor Pluggable (SFP) optical convertors (FC 2210) for distances up to 300 meters, or longwave optical convertors (FC 2220) for distances up to 10 kilometers (6.2 miles). If two SFPs are installed in a host-side miniHUB, a Fibre Channel loop connection is created between the FASTT miniHUB and the attached host server. On the FASTT900 drive interface, there are two drive-side miniHUBs standard, providing a redundant Fibre Channel loop pair to the dual ported Fibre Channel disk drives within the FASTT Expansion Units. Optionally, additional miniHUBs (FC 3507) can be added to the drive-side interface for a second pair of redundant Fibre Channel loops for drive attachment. If attaching more than eight of the maximum of 16 EXP700 Expansion Units, two additional drive-side miniHUBs are required. For optimum performance, it is recommended that all drive-side loops be utilized. The maximum number of host-side connections is eight and drive-side connections is four.

Refer to the Interoperability Matrix for specific details on supported operating system levels, prerequisites, and required attachments. The Interoperability Matrix is located on the Web at:

<http://www.ibm.com/storage/fast900>

FASTT Storage Manager Version 8.3 is included with the FASTT900. The FASTT Storage Manager provides a GUI for configuring, monitoring, and managing one or more FASTT Storage Servers. Included with the FASTT Storage Manager is a license for 16 Storage Partitions, allowing up to 16 different hosts or host clusters to attach to the FASTT900. Each Storage Partition provides protected data access for only the assigned host servers to the assigned FASTT logical volumes (or LUNs). An optional feature (FC 7664) provides a license activation key to increase the number of Storage Partitions up to 64. Each Storage Partition supports up to 32 logical volumes. With the 64 Storage Partition option, the FASTT900 supports up to 2048 logical volumes. The FASTT Storage Manager provides dynamic array management capabilities to expand logical volumes, add additional disk drives to array groups, and change the RAID segment size or RAID level of an array group without interrupting data availability on the FASTT.

Advanced Copy Services for Business Continuance

The Flashcopy option (FC 7301), available for the FASTT900, provides a license activation key to enable creating a point-in-time image of a logical volume within the FASTT. This feature can be combined with various backup or copy applications to provide data protection with nondisruptive data backups. The Remote Mirror option (FC 7302), available for the FASTT900, provides a license activation key to enable you to maintain a remote mirror copy of a logical volume on another FASTT Storage Server located at a disaster recovery site. This feature

provides a synchronous copy of your primary data across an extended distance, across a Fibre Channel fabric to a secondary site. Remote mirroring can be implemented in both directions between complementary FAST Storage Servers, primary data mirror from location to a secondary, and data on different logical volumes mirrored from the secondary site back to the primary. The Remote Mirror feature enables you to protect your critical data from catastrophic loss.

Limitations

The next section describe the limitations of the FAST900 Storage Server:

- ▶ FAST900 Storage Server is supported only in rack installations.
- ▶ FAST900 requires at least one FAST EXP700 or EXP500 Expansion Unit. Note that when configured with an EXP500 Expansion Unit, the drive-side FC loops will operate in 1 Gbps mode. For optimum performance, it is recommended that only EXP700 Expansion Units be configured with the FAST900.

Hardware and software requirements

For a list of hardware and software requirements, cross product requirements, prerequisites, and product dependencies, refer to the FAST900 Installation Guide and the following Web site:

<http://www.ibm.com/storage/fast900>

Devices supported

For a list of currently supported servers, operating systems, host bus adapters, clustering applications, SAN switches and directors, refer to the FAST900 Interoperability matrix available at the following Web site:

<http://www.ibm.com/storage/fast900>

Physical characteristics

The following describes the physical characteristics of the FAST900 storage server.

- ▶ Height: 174.5 mm (6.87 in.)
- ▶ Width: 482 mm (18.97 in.)
- ▶ Depth: 635 mm (24.0 in.)
- ▶ Weight: 43.9 kg (97 lb)

12.2.5 IBM TotalStorage FAStT700 Storage Server

IBM FAStT700 Storage Server delivers superior performance with 2 Gbps Fibre Channel technology. The FAStT700 is designed to offer investment protection with advanced functions and flexible features. It scales from 36 GB to over 32 TB to support growing storage requirements created by e-business applications. FAStT700 offers advanced replication services to support business continuance. The FAStT700 is an effective storage server for any enterprise seeking performance without borders.

Figure 12-5 shows a FAStT700 Model 1RU.




Figure 12-5 The FAStT700 Model 1RU

Highlights

The following list describes the highlights of the FAStT700 storage server:

- ▶ New 146.8 GB 10K RPM HDD expands capacity to 32.8 TB
- ▶ 2 Gbps high performance storage server
- ▶ RAID levels 0, 1, 3, 5, and 10
- ▶ FAStT Storage Manager Version 8.0 - Up to 64 storage partitions
- ▶ Up to 2048 logical drives, capable of nondisruptive expansion
- ▶ Optional functions - FlashCopy and Remote Mirror
- ▶ Higher scalability - Supports up to 220 FC HDDs with redundant loops
- ▶ Redundant, hot-swap components, and customer replaceable units (CRUs)
- ▶ 2 Gbps hot plug PCI-X host bus adapter
- ▶ Support for long-and short-wave FC environments with small formfactor pluggable (SFP) Gigabit Interface Converters (GBICs)
- ▶ Flexible SAN configurations supports pooling of disks and tape drives

- ▶ Remote data mirroring for up to 10 km (6.2 miles) for disaster recovery

Description

The IBM 1742 FAStT700 Storage Server Model 1RU expands the family of FAStT highly scalable Fibre Channel (FC) offerings from IBM. This third generation offering brings greater performance, advanced functionality, and expanded connectivity to Intel-based servers as part of a heterogeneous storage area networks (SANs).

- ▶ Compact 4 U rack-mountable unit features:
 - Dual 2 Gbps RAID controllers with 2 GB cache, 1 GB per controller
 - FAStT Storage Manager Version 8.0 - Up to 64 storage partitions, 32 logical drives per partition, for a total of 2048 logical drives
 - Two host and two drive-side FAStT Mini Hubs standard, four additional Mini Hubs optional
 - Up to eight host connection points. Dual redundant drive side loop capability with up to 110 drives per loop for a total of 220 drives
 - RAID levels 0, 1, 3, 5, and 10 supported
- ▶ TotalStorage FAStT FC-2 Host Bus Adapter
 - 64-bit, half-length, low profile, PCI-X host adapter
 - 2 Gbps technology - Autosenses 1 Gbps or 2 Gbps operation
 - Hot-plug and hot-add support without powering down the server
 - Supports point-to-point, FC-AL, FC-AL-2, and switched fabric topologies with FC SCSI and IP protocols
- ▶ Connection Options
 - Long-and Short-Wave® SFP GBICs
 - FAStT700 Mini Hub - Additional FAStT700 expansion capability
 - LC-LC Fibre Channel Cables -1, 5, and 25 M lengths
 - LC-SC Fibre Channel Adapter Cable - Connects small formfactor (LC) devices to current devices using standard connectors (SC)
- ▶ Short-and Long-Wave SFP GBICs

These SFP GBICs are used to connect cables into the FAStT700 Storage Server and other 2 Gbps Fibre Channel products. They provide short-and long-wave FC connections to other FAStT devices. They support the following installations:

 - Short-wave SFP GBIC - Distances up to 300 m (.18 miles)
 - Long-wave SFP GBIC - Distances of up to 10 km (6.2 miles)

Limitations

The next section describe the limitations of the FAStT700 Storage Server.

- ▶ TotalStorage FAStT700 Storage Server is supported only in rack installations.
- ▶ TotalStorage FAStT700 Storage Server supports optical cabling. Copper cabling is not supported.
- ▶ For optimum performance, no more than two TotalStorage FAStT700 Storage Servers should be attached to a single managed hub.
- ▶ FAStT500 feature Code Upgrades 7108 (FAStT Storage Manager Version 8.0), 7201 (Flashcopy Activation), and 7202 (Remote Mirror Activation) to the 3552-1RU are supported when the FAStT500 is attached to Intel-based platforms only.

Hardware requirements

The TotalStorage FAStT700 Storage Server is supported in two configurations:

- ▶ Connected directly to a TotalStorage FAStT FC-2 Host Bus Adapter through short-wave cabling
- ▶ Connected to an IBM SAN Fibre Channel Switch or SAN Fibre Channel Hub using short-wave cabling (the FC switch or hub must be interfaced either directly to a TotalStorage FAStT FC-2 Host Bus Adapter or through a second hub to a TotalStorage FAStT FC-2 Host Bus Adapter)

Software requirements

For a list of software requirements, cross product requirements, prerequisites, and product dependencies, refer to the following Web site:

<http://www.ibm.com/storage/fast700>

Rack requirements

The next section describes the rack requirements of the FAStT700 Storage Server.

- ▶ 19-in. EIA-310-D industry-standard rack such as the NetBAY42 SR, NetBAY25 SR, and the deeper NetBAY42 ER
- ▶ Mounts on standard EIA 19-in. rails (rails included with TotalStorage FAStT700 Storage Server)
- ▶ Height: 4 U (1 U = 1.75 in.) 177.8 mm (7.0 in.)
- ▶ Depth: 629.9 mm (24.8 in.) minimum

Physical characteristics

The following describes the physical characteristics of the FAStT700 storage server.

- ▶ Height: 74.5 mm (6.87 in.)
- ▶ Width: 444.5 mm (18.0 in.)
- ▶ Depth: 609.6 mm (24.0 in.)
- ▶ Weight: 38.6 kg (85 lb.)

12.2.6 IBM TotalStorage FAStT600 Storage Server

The FAStT600 is the latest addition to the FAStT family of products. The FAStT600 is an entry level, highly scalable 2 Gb Fibre Channel Storage Server offering best-in-class performance for an unbeatable solution in UNIX and Intel workgroups.

Figure 12-6 shows a FAStT600 Model 60U.




Figure 12-6 The FAStT600 Model 60U

Highlights

The following list describes the highlights of the FAStT600 storage server:

- ▶ Compact 3U rack-mount enclosure, offering a maximum internal storage capacity of over 2 TB, scalable to over 6 TB with EXP700 Expansion Units
- ▶ Redundant, hot-swappable components, helping to avoid single points of failure
- ▶ Support for RAID 0, 1, 3, 5, and 10, allowing for a choice of data protection based on fault-tolerance and performance requirements of individual user and application groups

- ▶ Storage-management software for configuring and continuously monitoring status
- ▶ Support for a wide range of industry-standard SAN switches and directors, enabling storage resources to be shared efficiently and scaled rapidly
- ▶ Supports high capacity and high performance 2 Gb FC HDDs for total capacity over 6 TB

Description

The FAST600 is designed to offer a highly available, high-performance, entry-level storage solution supporting a number of industry standards and can scale in capacity and connectivity within a wide range of SAN environments. FAST600 storage resources are designed to be shared efficiently and scaled rapidly to help address increased demands by users across a heterogeneous FC server environment.

The IBM TotalStorage FAST family of storage products are designed to offer a set of building blocks for Fibre Channel (FC)-based SANs, whether just starting or expanding an existing network. The FAST product family offers scalable storage capacity ranging from 36 GB to over 32 TB of capacity to support growing storage requirements. FAST Storage Servers are designed to be easy to install and manage using the FAST Storage Manager that is included with every FAST Storage Server. FAST products offer various controller models at different performance and price points to best address customer SAN needs.

Multiple FAST Storage Servers can be centrally managed from a common FAST management console from virtually anywhere within an existing LAN infrastructure. Use FAST advanced copy service functions, such as FlashCopy and Remote Mirror, to help address data protection and disaster recovery needs.

Key benefits:

- ▶ High availability: Dual processors and many redundant hardware components
- ▶ Scalability: From 36 GB to over 32 TB
- ▶ Performance: Excellent throughput to help match application needs
- ▶ Advanced functions to support business continuance: FlashCopy and Remote Mirror

The FAST600 utilizes 2 Gbps Fibre Channel technology for connections to host servers or switched SAN fabrics, and optionally to additional disk drive enclosures. On the FAST600 host interface, there are two host ports per controller, each providing two 2 Gbps FC attachment to each of the RAID controllers. You can attach to a switched fabric or directly to host servers using FC host bus adapters using either short-wave small form factor pluggable (SFP)

optical converters (FC 2210) for distances up to 300 meters, or long-wave SFP optical converters (FC 2220) for distances up to 10 km (6.2 miles). The FAStT600 supports and comes with authorization for up to 14 internal disk drive modules, supporting over 2 TB of storage capacity. Additional authorizations can be ordered for the FAStT600, allowing supported attachment of up to two FAStT EXP700 Expansion Units using optional EXP700 Attachment authorization features (FC 7360, or FC 7361 and FC 7362). Utilizing these features, up to 42 FAStT disk drive modules can be attached to the FAStT600. Individual drive module capacities can range from 36.4 GB to 146.8 GB.

Limitations

The next section describe the limitations of the FAStT600 Storage Server.

- ▶ The FAStT600 Storage Server is supported only in rack installations.
- ▶ With optional features, up to two FAStT EXP700 Expansion Units can be attached to the FAStT600 enclosure for a maximum of 42 disk drives.

Devices supported

For a list of currently supported servers, operating systems, host bus adapters, clustering applications, SAN switches and directors, refer to the FAStT600 Interoperability Matrix available at the following Web site:

<http://www.ibm.com/storage/fast600>

Hardware and software requirements

For a list of hardware and software requirements, cross-product requirements, prerequisites, product dependencies, and the FAStT600 Product Interoperability Matrix and other documentation, visit:

<http://www.ibm.com/storage/fast600>

Physical characteristics

The following describes the physical characteristics of the FAStT600 storage server.

- ▶ Height: 132.3 mm (5.21 in.)
- ▶ Width: 481.8 mm (18.97 in.)
- ▶ Depth: 597.4 mm (23.52 in.)
- ▶ Weight: 31.45 kg (69.34 lb)¹
- ▶ Weight: 45.87 kg (101.26 lbs)²

¹ Drive Ready Weight (without disk drive modules installed)

² Weight with 14 drives (fully configured)

12.2.7 IBM TotalStorage FAStT200 Storage Server

The 3542 FAStT200 is designed for workgroup and departmental servers that require an external storage solution. The single controller model provides a cost-effective solution, while the FAStT200 High Availability (HA) model features a fully redundant configuration with dual-active controllers. As your storage requirements grow, you can easily expand storage capacity by adding IBM FAStT EXP500 or EXP700 Expansion Units.

Figure 12-7 shows a FAStT200 Model 1RU.



Figure 12-7 The FAStT200 Model 1RU

Highlights

The following list describes the highlights of the FAStT200 storage server:

- ▶ Four compact models to suit customer needs.
- ▶ High-availability technology.
- ▶ High-performance 1 Gbps Fibre Channel technology.
- ▶ Scalability.
- ▶ Data storage up to 10 km (6.2 miles) away: Provides additional protection from catastrophic occurrences.
- ▶ Support for long-wave and short-wave Fibre Channel environments.
- ▶ Supports up to 1.46 TB storage capacity per unit utilizing 10 K RPM 146.8 GB drives.
- ▶ 15 K RPM 18.2 GB, 36.4 GB, and 73.4 GB Fibre Channel drive options.
- ▶ Provides a wide range of data protection options with RAID levels 0, 1, 3, 5, and 10. Fully integrates Fibre Channel technology, from host attachment to disk drives.

Description

If you need a low-cost, space-efficient storage solution for the exploding Internet markets, then you will be interested in the IBM FAStT200 Storage Server. This attractive disk storage server is ideal for small to medium businesses that require a low-cost, full-fibre solution that is compact and Storage Area Network (SAN)-ready for Intel-based servers.

The FAStT200 Storage Server is available in two versions: single controller FAStT200 (Models 1RX or 1RU) or dual controller FAStT200 HA (Models 2RX or 2RU). The single controller version will provide low-entry price point, whereas the dual controller version should be implemented for high-availability needs. All configurations deliver up to 734 GB of storage capacity in a single 3U-high, rack-mounted package.

The FAStT200 Storage Server supports up to two IBM FAStT EXP500 Expansion Units. Each EXP500 Expansion Unit houses up to 10 Fibre Channel hot-swap disk drives to provide a storage capacity of up to 1.46 TB per unit. It supports 10,000 rpm 18, 36, 73, or 146 GB FC HDDs, or the 15,000 RPM 18 GB, 36 GB, and 73 GB FC HDDs. The EXP500 is a fully-redundant, highly available unit with hot-swap dual power supplies and fans.

The IBM FAStT200 Storage Server is available in four models: 1RU, 1RX, 2RU, and 2RX.

Key features of models 1RU and 1RX

The following describes the key features of the IBM FAStT200 Storage Server models 1RU and 1RX. These single controller models are a cost-effective solution for applications where a high degree of redundancy is not an immediate requirement.

- ▶ Includes 1 hot-plug RAID controller, 10 disk drive bay, redundant fans, and power supplies
- ▶ Provides single Fibre Channel Arbitrated Loop (FC-AL) host connectivity and 1 drive loop
- ▶ Supports hot-swap disk drives with capacities of 18 GB, 36 GB, 73 GB, and 146 GB (all 10 K RPM)
- ▶ Optional 18 GB, 36 GB, and 73 GB HDDs with 15 K RPM
- ▶ Supports RAID levels 0, 1, 3, 5, and 10
- ▶ Contains Netfinity FAStT Storage Manager Version 7.02: Supports up to four storage partitions
- ▶ Supports both short and long wave Fibre Channel connections

- ▶ Hot-plug single RAID controller with:
 - 100 MHz Intel 80960RN microprocessor with 62-bit, 33 MHz primary and secondary PCI buses
 - 128 MB of ECC, battery backup cache
 - One expansion port GBIC connector: Allows connections through an optional GBIC to a Netfinity FAStT EXP500
 - 10/100 Mbps Ethernet connector
 - RS-232 service support

Key features of models 2RU and 2RX

The following describes the key features of the IBM FAStT200 Storage Server models 2RU and 2RX. These dual controller models provide redundancy and higher availability to 1RU/1RX models. These models offer similar features to 1RU/1RX models with an additional controller unit.

- ▶ Includes two hot-plug RAID controllers, 10 disk drive bays, redundant fans, and power supplies
- ▶ Provides dual FC-AL host connectivity and 2 drive loop: one host and one drive loop per controller
- ▶ Supports hot-swap disk drives with capacities of 18 GB, 36 GB, 73 GB, and 146 GB (all 10 K RPM)
- ▶ Optional 18 GB, 36 GB, and 73 GB HDDs with 15K RPM
- ▶ Supports RAID levels 0, 1, 3, 5, and 10
- ▶ Contains Netfinity FAStT Storage Manager Version 7.02: Supports up to four storage partitions
- ▶ Supports both short and long wave Fibre Channel connections
- ▶ Performance optimized for up to 60 disk drives, using optional FAStT EXP 500 Expansion units
- ▶ Dual hot-plug RAID controller:
 - 100 MHz Intel 80960RN microprocessor with 62-bit, 33 MHz primary and secondary PCI buses
 - 256 MB of ECC, battery backup cache
 - Two expansion port GBIC connector: Allows connections through an optional GBIC to a Netfinity FAStT EXP500
 - 10/100 Mbps Ethernet connector
 - RS-232 service support

- ▶ In addition to direct attach (using FAStT Host Adapter), all four models support connection to the host servers using Fibre Channel SANs (IBM Switches and Hubs). The FAStT200 Storage Servers are shipped without HDD, GBIC and cable options to provide flexibility to satisfy various storage requirements. The FAStT200 Storage Server can be upgraded to a redundant unit by installing a FAStT200 Failsafe RAID controller. The FAStT200 Storage Servers are supported in two configurations:
 - Connected directly to a Netfinity FAStT Host Adapter through short-wave cabling
 - Connected indirectly through an IBM Fibre Channel Switch or an IBM Fibre Channel SAN Hub using either short or long-wave cabling

For more information about servers supported with this adapter, refer to the Web site at:

<http://www.pc.ibm.com/us/compat/controllers/matrix>

Hardware and software requirements

For a list of hardware and software requirements, cross product requirements, prerequisites, and product dependencies, refer to the following Web site:

<http://www.ibm.com/storage/FAStT200>

Physical characteristics

The following describes the physical characteristics of the FAStT200 storage server.

- ▶ Width: 480 mm (18.9 in.)
- ▶ Depth: 575 mm (22.63 in.)
- ▶ Height: 131.8 mm (5.2 in.)
- ▶ Weight: 25.4 kg (56 lb.)

12.3 IBM TotalStorage Tape drive products

This section describes the various IBM TotalStorage tape products that can be attached to the pSeries machines. The following tape drive products are discussed:

- ▶ 7205 IBM External Digital Linear Tape Drive
- ▶ 7206 IBM External 4 mm Tape Drive
- ▶ 7207 IBM 4 GB External SLR5 QIC Tape Drive Model 122

12.3.1 7205 IBM External Digital Linear Tape Drive

Figure 12-8 shows the 7205 External Digital Linear Tape Drive.



Figure 12-8 7205 External Digital Linear Tape Drive

The IBM 7205 is an external, digital, linear, stand-alone tape drive.

The External Digital Linear Tape Drive 7205 is available in two models for SCSI Ultra2 LVD (68-pin) attachment: Models 440 and 550.

Model highlights of the 7205-440

The following are the highlights of the 7205-440:

- ▶ Increased tape capacity of 40 GB native per cartridge (80 GB with 2:1 compression) using DLTtape IV
- ▶ Improved data transfer rate up to 6 MB per second (uncompressed), up to 12 MB per second with 2:1 compression
 - Ultra2/Wide Low Voltage Differential (LVD) SCSI interface

Model highlights 7205-550

The following are the highlights of the 7205-550:

- ▶ Provides four times the capacity of the IBM TotalStorage 7205 Model 440 D
- ▶ Features a sustained data rate of 16 MBps native and 32 MBps compressed
- ▶ Offers up to 160 GB (320 GB compressed) capacity per cartridge
- ▶ Offers compatibility with 7205 Model 440 Tape Drives used on IBM @server pSeries and IBM RS/6000 servers
- ▶ To help protect the data availability of tapes written by older generation DLT drives, the 7205 Model 550 is designed to be capable of reading tapes written on other systems that use SDLT320, DLT7000, or DLT8000 tape, enabling cost-effective backup, restore, and archiving functions. As DLT8000 tape technology nears the end of its evolution, the new SDLT320 standard, through

its backward compatibility with older DLT drives, can provide a migration path to greater tape storage capacity at a price point similar to DLT drives.

Table 12-5 summarizes the key features of the 7205 IBM External Digital Linear Tape Drive.

Table 12-5 IBM 7205 External SDLT320 Tape Drive at a glance

Specification	Detail
Model Number	440, 550
Capacity per cartridge	440: Up to 80 GB compressed, 40 GB native 550: Up to 320 GB compressed, 160 GB native
Data transfer rate	440: 6 MB per second (uncompressed), up to 12 MB per second with 2:1 compression 550: 16 MBps native, 32 MBps compressed ¹
Interface	SCSI Ultra2 LVD (68-pin)
Operating systems	IBM AIX Version 4.3.3, AIX 5L Version 5.1, or AIX 5L Version 5.2 or later
Media	S-DLT Data Cartridge (160 GB) (P/N 35L1119) S-DLT Test Cartridge (P/N 35L1120) S-DLT Cleaning Cartridge (P/N 19P4357)
Dimensions	122 mm H x 250 mm W x 290 mm D (4.8 in. x 9.8 in. x 11.5 in.)
Weight	6.3 kg (13.8 lbs)
Warranty	24x7, one year, IBM on-site repair

12.3.2 7206 IBM External 4 mm Tape Drive

Figure 12-9 shows the 7206 External 4 mm Tape Drive.



Figure 12-9 7206 External 4 mm Tape Drive

The IBM 7206 External Tape Drive series is comprised of stand-alone, SCSI, VXA-2 or 4 mm (DDS-4) tape drives. The available models are the following:

- ▶ 7206-220 DDS-4
- ▶ 7206-VX2

7206-220 DDS-4

The following are the highlights of the 7206-220:

- ▶ Up to 20 GB/40 GB data storage capacity
- ▶ Up to 3 MBps/6 MBps data transfer rate

7206-VX2

The following are the highlights of the 7206-VX2:

- ▶ Up to 80 GB/160 GB data storage per cartridge
- ▶ Up to 6.0 MBps data transfer rate

Table 12-6 summarizes the key features of the 7206-220.

Table 12-6 IBM 7206-220 Model DDS-4 at a glance

Specification	Detail
Capacity per cartridge	40 GB compressed, 20 GB uncompressed
Data transfer rate	6 MBps compressed, 3 MBps uncompressed
Interface	LVD/SE W Ultra SCSI-2
pSeries Hardware requirements	SCSI-2 Fast/Wide PCI-Bus Adapter (FC 2408), SCSI-2 Fast/Wide Adapter/A (FC 2415), PCI Dual Channel Ultra 2 SCSI Adapter (FC 6205), PCI Single-Ended Ultra SCSI Adapter (FC 6206), PCI SCSI-2 Fast/Wide Single-Ended Adapter (FC 6208)
Operating Systems	AIX Version 4.2, 4.3, or later
Media	4 mm DDS-4 Data cartridge (P/N 59H4458) 4 mm DDS-4 Test Cartridge (P/N 59H4457) 4 mm DDS-4 Cleaning Cartridge (P/N 21F8763)
Warranty	One-year limited warranty
Dimensions	2.2 in. H x 9.8 in. W x 10.8 in. D (55 mm x 250 mm x 275 mm)

Table 12-7 on page 483 summarizes the key features of the 7206-VX2.

Table 12-7 IBM 7206-VX2 Model VX2 at a glance

Specification	Detail
Capacity per cartridge	160 GB compressed, 80 GB uncompressed
Data transfer rate	6 MBps
Interface	Ultra2 SCSI LVD (68-pin)
pSeries hardware requirements	<ul style="list-style-type: none"> ▶ FC 6203 PCI Dual Channel Ultra3 SCSI Adapter ▶ FC 6205 PCI Ultra-2 SCSI LVD (low voltage differential or single-ended) 64 bit Adapter ▶ FC 6206 PCI Single-Ended Ultra SCSI Adapter ▶ Integrated Ultra2 SCSI Adapter (LVD with VHDCI connector) ▶ Integrated Ultra3 SCSI Adapter (LVD with VHDCI connector) <p>Please check with the sales representatives for specific host adapters supported by model and detailed attachment and configuration information.</p>
Operating systems	AIX Version 4.3.0, 5.1.0, or later.
Media	<ul style="list-style-type: none"> ▶ VXA-2 19P4876 230 m VXA-2 Data Cartridge (80 GB) 19P4877 170 m VXA-2 Data Cartridge (59 GB) 19P4878 62 m VXA-2 Data Cartridge (20 GB) 19P4879 VXA-2 Test Cartridge 19P4880 VXA-2 Cleaning Cartridge
Warranty	24x7, one year, IBM on-site repair
Dimensions	<ul style="list-style-type: none"> ▶ Width: 250 mm (9.8 in.) ▶ Depth: 275 mm (10.8 in.) ▶ Height: 55 mm (2.2 in.) ▶ Weight: 3.4 kg (7.5 lbs)

12.3.3 7207-122 IBM 4 GB external SLR5 QIC tape drive Model 122

Figure 12-10 shows the 7207-122 tape drive.



Figure 12-10 7207-122 IBM 4 GB External SLR5 QIC Tape Drive

The IBM 7207 Model 122 is an excellent solution if you use quarter-inch cartridge (QIC) tape or require a low-cost entry solution for tape backup. This economical entry-level save and restore device for medium-capacity fixed disk drive data and applications uses the popular QIC data cartridge for data exchange and software distribution.

Table 12-8 summarizes the key features of the 7207-122.

Table 12-8 IBM 7207-122 4 GB SLR5 QIX Tape drive at a glance

Specification	Detail
Tape drive type	QIC (1/4-in.)
Capacity per cartridge	8 GB compressed, 4 GB uncompressed
Max. data transfer rate	380 KB/s (up to 760 KB/s with compression)
Interface	SCSI-2

Specification	Detail
pSeries hardware requirements	<ul style="list-style-type: none"> ▶ FC 2415 SCSI-2 Fast/Wide PCI-Bus Adaptor ▶ FC 2410 SCSI-2 High-Performance External I/O Controller ▶ FC 6206 PCI Single-Ended Ultra SCSI Adaptor ▶ FC 6208 PCI SCSI-2 Fast/Wide Single-Ended Adapter ▶ Integrated SCSI-2 port on MCA systems ▶ Integrated SCSI-2 Fast/Wide port on MCA systems ▶ Integrated SCSI-2 Fast/Wide single-ended port on PCI systems ▶ Integrated Ultra SCSI-2 Fast/Wide single-ended port on PCI systems
Operating systems	<ul style="list-style-type: none"> ▶ IBM AIX Version 4.1.5 or later ▶ IBM AIX Version 4.2.0 and 4.2.1 or later ▶ IBM AIX Version 4.3.0 or later
Media	4 GB data cartridge (P/N 59H3660)
Warranty	IBM Onsite repair, 1 year limited

12.4 IBM Tape Libraries/Autoloaders

This section covers the following tape library/autoloader products available from IBM:

- ▶ 3590 IBM TotalStorage Enterprise Tape System
- ▶ 3494 IBM TotalStorage Enterprise Automated Tape Library
- ▶ 7332 IBM 4 mm DDS-2, DDS-3, and DDS-4 Tape Cartridge Autoloader

12.4.1 3590 IBM TotalStorage Enterprise Tape System

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.

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 IBM TotalStorage Enterprise Tape System 3590 provides high levels of capacity, performance, and reliability for IBM @server iSeries (AS/400), IBM @server pSeries (RS/6000), IBM @server xSeries, IBM @server zSeries (S/390), and a variety of non-IBM UNIX and Microsoft Windows systems. It offers significant enhancements over previous IBM 3480, 3490, or 3490E tape drives and demonstrates IBM's continued leadership in storage solutions. The Enterprise Tape System 3590 includes tape drives and controllers that can be flexibly configured in a broad range of solutions that includes stand-alone frames or racks, or in automated library solutions to meet a broad set of customer requirements. The 3590 Tape Drives can be attached to a SCSI host interface or Fibre Channel interface of selected IBM and non-IBM midrange and open systems. The tape drives can also attach to ESCON and FICON channels on host systems through a 3590 controller:

The Magstar 3590 is available in six models for Ultra SCSI host and Fibre Channel attachment.

- ▶ The 3590 Models B11, E11, and H11 are rack-mounted and incorporate a 10-cartridge ACF for high-capacity unattended operation. The Models B11, E11, and H11 can be converted to Models B1A, E1A, or H1A.
- ▶ The 3590 Models B1A, E1A, and H1A have no ACF and are designed to be incorporated into the 3494 Enterprise Tape Library.
- ▶ The 3590 Model C12 Silo Compatible Frame (with one to four Model B1A, E1A, or H1A tape drives) provides attachment to the StorageTek 4410 and 9310 ACS.

Table 12-9 describes the 3590 tape drive models.

Table 12-9 3590 tape drive models at a glance

Model	Description	Attaches to	Interface
B11	Tape drive with ACF	Model A14 frame or 9309, 7014, or 7015	16-bit, fast-and-wide, Ultra SCSI differential.
B1A	Tape drive without ACF	3590 Model A14, C12, C14; 3494 library	16-bit, fast-and-wide, Ultra SCSI differential.

Model	Description	Attaches to	Interface
E11	Tape drive with ACF	Model A14 frame or 9309, 7014, or 7015	16-bit, fast-and-wide, Ultra SCSI differential interface or a dual ported Fibre Channel attachment.
E1A	Tape drive without ACF	3590 Model C12, C14; 3494 library	16-bit, fast-and-wide, Ultra SCSI differential interface or dual ported Fibre Channel attachment.
H11	Tape drive with ACF	Model A14 frame, or 7014 or 7015	16-bit, fast-and-wide, Ultra SCSI differential interface or a dual ported Fibre Channel attachment.
H1A	Tape drive without ACF	3590 Model C12, C14 frame; 3494 library	16-bit, fast-and-wide, Ultra SCSI differential interface or a dual ported Fibre Channel attachment.
C12	Silo Compatible Frame	StorageTek ACS (4400 or 9310)	One to four IBM TotalStorage Enterprise Tape Drives 3590 Model H1A, E1A, or Model B1A can be installed in the Model C12 frame. The 3590 Tape Drives can be attached to the SCSI or Fibre Channel adapter.

12.4.2 3494 IBM TotalStorage Enterprise Automated 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/H1A and 3490E (withdrawn from marketing) tape drives within the same library, but in separate frames.

Highlights

The following are the highlights of the 3494 Tape Library:

- ▶ Provides cost-effective, reliable, and space-efficient tape automation for a variety of environments.
- ▶ Uses both Magstar 3590 Model B1A/E1A/H1A and 3490E tape drives (withdrawn).
- ▶ 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 with 3:1 compression).
- ▶ Can mount/demount up to 610 cartridges per hour with dual active accessors.
- ▶ 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.

Components of the 3494

Magstar 3494 library components consists of a control unit frame, a drive unit frame, and a storage unit frame.

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: Withdrawn from marketing. One or two 3490E Model F1A tape drives.
- ▶ Model L12: Zero to two SCSI-attached Magstar 3590 tape drives.

- ▶ Model L14: Withdrawn from marketing. Zero to two Magstar 3590 tape drives and an ESCON control unit (3590-A50).

Drive unit frame

Provides the capability of adding tape drives and up to 400 cartridge storage cells.

Configurations include:

- ▶ Model D10: Withdrawn from marketing. 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).
- ▶ Model B10: Zero to four FICON channels, four Extended Performance ESCON Channels, up to eight SCSI bus attachments.
- ▶ Model B20: Zero to eight FICON channels, up to 16 Extended Performance ESCON Channels, up to eight SCSI bus attachments.
- ▶ Model HA1: Operates in stand-by mode to provide a redundant library manager and accessor for improved availability. With the Dual Active Accessor (DAA) feature active on the IBM TotalStorage Enterprise Tape Library Base Frame Model L10, L12, or L14, both accessors can operate simultaneously to increase mount performance of the library.
- ▶ 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.

Storage unit 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 (withdrawn from marketing).
- ▶ 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.

Table 12-10 summarizes the key features of the 3494 Tape Library.

Table 12-10 IBM TotalStorage 3494 Tape Library at a glance

Specification	Detail
Max. control unit frames 3494-L12	oNe
Library Manager communications	LAN or RS-232 connection or through the ESCON and FICON or parallel channel controller
Data channel attachments	SCSI, Fiber Channel, parallel channel, ESCON, or FICON
Max. additional drive unit frames	15
Max. tape drives with SCSI or Fibre Channel attachments per L12	Zero to two 3590 Model H1A, E1A, or B1A
3494-L12 cartridge capacity	160 to 240
3494-L12 flexible growth	<ul style="list-style-type: none"> ▶ One to two SCSI drives ▶ One to 16 frames with up to 6,240 cartridges ▶ Technology supported: 3590 H1A, E1A, or B1A
3494-D12 drive unit frame capacity	<ul style="list-style-type: none"> ▶ Cartridge capacity/Flexible growth ▶ 400/0 drives ▶ 335/one to two drives ▶ 290/three to four drives ▶ 250/five to six drives ▶ Technology supported: 3590 H1A, E1A, or B1A
Hardware requirements	<ul style="list-style-type: none"> ▶ 3590 tape drives supported within the Model L12: <ul style="list-style-type: none"> – SCSI-attached: H1A, E1A, or B1A, – Fibre-Channel attached: H1A or E1A <p>The Model L12 may be attached to any of the following host processors for library commands: RS/6000 models, including SP models, that support SCSI-2 Differential adapters (FC 2409, 2412, 2416, 2420, 6207, and 6209), System/390 ESCON Channel Emulator adapter (FC 2754), and System/370™ Channel Emulator/A adapter (FC 2759)</p>

Specification	Detail
AIX operating system	<p>An enhanced version of the Open Systems Device Driver (FC 9200) will be required to support a configuration containing the Model HA1 Frames.</p> <ul style="list-style-type: none"> ▶ AIX Version 4.1.1, and later, for 3490E support ▶ RS/6000 RPQ 8A1016 - ESCON Channel Tape Attachment ▶ 3494 specify feature#9200 (Open System Device Drivers) ▶ Tivoli® Storage Manager for AIX Version 2.1 plus PTFs, or Tivoli Storage Manager Version 3 ▶ AIX Version 4.3.3, and later, for 3590 Fibre Channel attached support ▶ AIX Version 4.1.5, and later, for 3590 SCSI support ▶ 3494 specify feature #9200 (Open System Device Drivers) ▶ Tivoli Storage Manager for AIX Version 2.1 plus PTFs, or Tivoli Storage Manager Version 3

12.5 LTO Ultrium tape drives, libraries, and autoloaders

This section provides a positioning guide and discusses the following LTO products:

- ▶ 3580 IBM Ultrium Tape Drive
- ▶ 3581 IBM Ultrium Tape Autoloader
- ▶ 3582 IBM TotalStorage Ultrium Tape Library Model L23
- ▶ 3583 IBM Ultrium Scalable Tape Library
- ▶ 3584 IBM UltraScalable Tape Library

12.5.1 LTO product positioning notes

For capacity requirements greater than 28.8 TB (compressed), the IBM TotalStorage UltraScalable Tape Library 3584 should be considered. The UltraScalable Tape Library 3584 can match system capacity and performance requirements from 28 TB to 496 TB (56 TB to 992 TB with 2:1 compression) using up to 72 IBM Ultrium Tape Drives in up to six 3584 library frames.

For capacity requirements in the 4.8 TB to 9.6 TB (compressed) range with an entry-level automation need, the new IBM TotalStorage Ultrium Tape Library

3582 should be considered with its one or two IBM LTO Ultrium 2 Tape Drives and 23-cartridge capacity.

For capacity requirements less than 1.4 TB (compressed), the IBM TotalStorage Ultrium Tape Autoloader 3581 and the IBM TotalStorage Ultrium Tape Drive 3580 should be considered.

For mission-critical data protection, optimized for enterprise multi-mode and host attachment, and high-cycle and start/stop intensive tape applications, consider the proven IBM TotalStorage Enterprise Tape Drive 3590 or the IBM TotalStorage Enterprise Automated Tape Library 3494.

12.5.2 3580 IBM Ultrium Tape Drive

Figure 12-11 shows the IBM 3580 Ultrium Tape Drive.



Figure 12-11 IBM 3580 Ultrium Tape Drive

The IBM 3580 is the Ultrium Tape Drive with industry-leading software applications that offer high capacity and performance. Its technology is designed for the heavy demands of backup tape storage.

3580 is available in four models for Ultra2/Wide SCSI LVD or Ultra/Wide SCSI HVD interface using Ultrium 1 drives and two models using Ultrium 2 drives.

Table 12-11 on page 493 summarizes the key features of the 3580 models using Ultrium 1 drives.

Table 12-11 IBM Ultrium Tape Drive at a glance using Ultrium 1 drives

Specification	Detail
Model number	L11—LVD attach; H11—HVD attach; one year on-site warranty L13—LVD attach; H13—HVD attach; three year customer exchange warranty
Tape drive type	Ultrium 1
Number of tape drives	1
Number of tape cartridges	1
Capacity per cartridge	Up to 200 GB compressed; 100 GB native
Interface	L11 and L13 features Ultra2/Wide SCSI LVD interface L13 and L17 features Ultra/Wide SCSI HVD interface
pSeries hardware requirements	<ul style="list-style-type: none"> ▶ FC 6204 (H11/H13) PCI Universal Differential Ultra SCSI (HVD) (requires AIX Version 4.3.3, or later) ▶ FC 6203 (L11/L13), PCI Dual Channel Ultra3 SCSI Adapter (LVD/SE with VHDCI connector) ▶ FC 6207 (H11/H13) PCI Differential Ultra SCSI Adapter (HVD) ▶ FC 6205 (L11/L13) PCI Dual Channel Ultra2 SCSI Adapter (LVD/SE with VHDCI connector) ▶ (L11/L13) Integrated Ultra2 SCSI Adapter (LVD with VHDCI connector) ▶ (L11/L13) Integrated Ultra3 SCSI Adapter (LVD with VHDCI connector) <p>Check with the sales representatives for specific host adapters supported by model, detailed attachment and configuration information.</p>
Operating systems	<p>The following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 4 Release 4 or later ▶ AIX Version 4.3.2 or later ▶ Sun Solaris Version 2.6, Solaris Version 7, Solaris Version 8 ▶ Microsoft Windows NT Version 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 ▶ HP-UX 11.0

Specification	Detail
Media	<ul style="list-style-type: none"> ▶ IBM Ultrium ▶ Data Cartridge Part number 08L9120 ▶ Cleaning Cartridge Part number 08L9124
Dimensions	.75" H x 6.74" W x 13.11" D (14.6 cm x 17.1 cm x 33.3 cm)
Weight	14.3 lb (6.6 kg)
Warranty	<ul style="list-style-type: none"> ▶ L11; H11 - One year; IBM On-site Exchange (IOE) ▶ L13; H13 - Three year, Customer Element Exchange (CEE)

Table 12-12 summarizes the key features of the 3580 models using Ultrium 2 drives.

Table 12-12 3580 IBM Ultrium Tape Drive at a glance using Ultrium 2 drives

Specification	Detail
Model number	L23—LVD attach; H23—HVD attach
Tape drive type	IBM LTO Ultrium 2
Number of tape drives	1
Number of tape cartridges	1
Capacity per cartridge	Up to 400 GB compressed; 200 GB native
Interface	<ul style="list-style-type: none"> ▶ L23 features Ultra SCSI LVD interface ▶ H23 features Ultra SCSI HVD interface
pSeries hardware requirements	A current list of supported open system configurations is available from the following Web site: http://www.ibm.com/storage/lto

Specification	Detail
Operating systems	<p>The Ultrium 2 Tape Drive 3580 is supported on the following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 5 Release 1 or later ▶ AIX Version 4.3.3 or later ▶ Sun Solaris 7, 8, or 9 ▶ Microsoft Windows NT Version 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 (build 2195 or greater) ▶ HP-UX 11.0 and HP-UX 11i (64-bit) ▶ Linux distributions: Linux 7.2 32-bit and 64-bit Kernels, Linux Red Hat 7.3, Red Hat Advanced Server 2.1, SuSE Linux Enterprise Server 7 Update
Media	<ul style="list-style-type: none"> ▶ IBM Ultrium 2 ▶ IBM TotalStorage LTO Ultrium 200 GB Data Cartridge (PN 08L9870) ▶ IBM TotalStorage LTO Cleaning Cartridge (PN 35L2086)
Dimensions	5.75" H x 6.74" W x 13.11" D (14.6 cm x 17.1 cm x 33.3 cm)
Weight	14.3 lb (6.6 kg)
Warranty	Three year, IBM 3580 media warranted separately

12.5.3 3581 IBM Ultrium Tape Autoloader

Figure 12-12 shows the 3581 Ultrium Tape Autoloader.



Figure 12-12 IBM 3581 Ultrium Tape Autoloader

The 3581 Ultrium Tape Autoloader is an external, stand-alone, or rack-mounted autoloader that incorporates an IBM Ultrium Tape Drive.

The 3581 Ultrium is available in four models for Ultra2/Wide SCSI LVD or Ultra/Wide SCSI HVD interface.

Highlights

The following are the highlights of the 3581 Ultrium Tape Drive:

- ▶ Tape capacity up to 100 GB native per cartridge (200 GB with 2:1 compression)
- ▶ Seven storage slots, an optional bar code reader, and front-panel LCD display
- ▶ Data transfer rate up to 15 MB per second (uncompressed), up to 30 MB per second with 2:1 compression
- ▶ Optional bar code reader
- ▶ External, self-powered packaging
- ▶ Front-panel LCD displays Support for a wide variety of backup/restore software
- ▶ Adherence to Linear Tape-Open (LTO) specifications

Table 12-13 summarizes the key features of the IBM 3581 Ultrium Tape Drive.

Table 12-13 IBM 3581 Ultrium Tape Drive at a glance

Specification	Detail
Model number	<ul style="list-style-type: none"> ▶ L17—LVD attach; H17—HVD attach; one year on-site warranty ▶ L13—LVD attach; H13—HVD attach; three year customer exchange warranty
Tape drive type	Ultrium
Number of tape drives	1
Number of tape cartridges	7
Capacity	<ul style="list-style-type: none"> ▶ Cartridge capacity: Up to 200 GB per cartridge compressed; 100 GB native ▶ Autoloader capacity: seven tape cartridges, providing a media capacity of up to 1.4 TB with 2:1 compression; 700 GB native
Interface	<ul style="list-style-type: none"> ▶ L13 and L17 features Ultra2/Wide SCSI LVD interface ▶ H13 and H17 features Ultra/Wide SCSI HVD interface
pSeries hardware requirements	<ul style="list-style-type: none"> ▶ FC 6203 (L13/L17), PCI Dual Channel Ultra3 SCSI Adapter (LVD/SE with VHDCI connector) ▶ FC 6205 (L13/L17), PCI Dual Channel Ultra2 SCSI Adapter (LVD/SE with VHDCI connector) ▶ FC 6204(H13/H17), PCI Universal Differential Ultra SCSI (HVD) (requires AIX Version 4.3.3, or later) ▶ FC 6207 (H13/H17), PCI Differential Ultra SCSI Adapter (HVD) ▶ (L13/L17) Integrated Ultra2 SCSI Adapter (LVD with VHDCI connector) (requires AIX Version 4.3.3, or later). ▶ (L13) integrated Ultra3 SCSI Adapter (LVD with VHDCI connector) <p>Check with the sales representatives for specific host adapters supported by model, detailed attachment and configuration information.</p>

Specification	Detail
Operating systems	Native device driver support is available for AIX Version 4.3.2 or 4.3.3; OS/400 Version 4 Release 4; Windows NT Version 4.0 with Service Pack 6.0; and Windows 2000; Sun Solaris 2.6, 7, or 8; HP-UX Version 11.0; and Red Hat Linux 7.0.
Media	<ul style="list-style-type: none"> ▶ IBM Ultrium Data Cartridge Part number 08L9120 ▶ Cleaning Cartridge Part number 08L9124
Dimensions	7.48 in. H x 8.62 in. W x 22.87 in. D (19.0 cm x 21.9 cm x 58.1 cm)
Weight	28.7 lb (13.0 kg)
Warranty	<ul style="list-style-type: none"> ▶ L17; H17 - One year; IBM Ultrium media warranted separately ▶ L13; H13 - Three year, IBM Ultrium media warranted separately

12.5.4 3582 IBM TotalStorage Ultrium Tape Library Model L23

Figure 12-13 shows the Ultrium Tape Library 3582 Model L23.



Figure 12-13 IBM TotalStorage Ultrium Tape Library 3582 Model L23

The IBM 3582 is the Ultrium Tape Library that incorporates high-performance IBM TotalStorage LTO Ultrium 2 Tape Drives for the midrange open systems environment.

The IBM 3582 Model L23 is an entry level Ultrium Tape Library that can accommodate one or two Ultrium 2 Tape drives and comes standard with a

one-cartridge I/O station and 23 data cartridge slots, giving a native library capacity of 4.8 TB uncompressed data storage (9.6 TB with 2:1 compression).

Highlights

The following are the highlights of the 3582 Ultrium Tape Library Model L23:

- ▶ Use of LTO Ultrium 2 Tape Drives with a drive data rate of up to 70 MBps (2:1 data compression)
- ▶ Up to two SCSI or native switched fabric Fibre Channel drive attachments
- ▶ 4U Height Tape Library; Stand-alone Desktop or Rack Mount
- ▶ Multi-Path Architecture/Logical Library Partitioning
- ▶ Two Logical Library Partitions
- ▶ One application per logical library partition
- ▶ Remote Management Unit
- ▶ Required for Fibre Channel (optional for SCSI)
- ▶ Barcode Scanner
- ▶ Control Path Failover for AIX Option
- ▶ Enables automatic failover to redundant HBA, SAN, or library control path when a failure is detected by the device driver software
- ▶ Implemented through library/drive firmware and IBM device driver updates
 - Available initially for Fibre Channel attachment

Table 12-14 summarizes the key features of the IBM TotalStorage Ultrium Tape Library 3582 Model L23.

Table 12-14 IBM TotalStorage Ultrium Tape Library 3582 Model L23 at a glance

Specification	Detail
Model number	L23 Base Library with 23 cartridges, 1-I/O slot.
Tape drive type	IBM LTO Ultrium 2.
Number of tape drives	Up to 2.
Number of tape cartridges	24.
Capacity per cartridge	<ul style="list-style-type: none"> ▶ Up to 400 GB per cartridge compressed; 200 GB native ▶ Up to 9.6 TB per library compressed; 4.8 TB native
pSeries hardware requirements	Native switched fabric 2 Gb Fibre Channel, LVD Ultra 160 SCSI and HVD Ultra SCSI interfaces.

Specification	Detail
Operating systems	<p>Supported on the following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 5 Release 1, or later ▶ AIX Version 4.3.3, or AIX 5L Version 5.1 or Version 5.2 or later ▶ Sun Solaris 7, 8, or 9 ▶ Microsoft Windows NT 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 (build 2195, or greater) ▶ Microsoft Windows 2003 (build 3790, or greater) ▶ HP-UX 11.00 and 11i (64-bit) ▶ Linux distributions: Red Hat Advanced Server 2.1, SuSE Linux Enterprise Server Update <p>For a current list of host software versions and release levels that support the 3582, refer to the following Web site: http://www.ibm.com/storage/lto Tivoli Storage Manager and other industry-leading compatible software offerings provides storage and tape management software for the 3582.</p>
Media	<ul style="list-style-type: none"> ▶ IBM TotalStorage LTO Ultrium 200 GB Data Cartridge (P/N 08L9870) ▶ IBM TotalStorage LTO Universal Cleaning Cartridge (P/N 35L2086)
Dimensions	7.7" H x 17.9" W x 25.8" D (19.4 cm x 45.5 cm x 65.4 cm)
Weight	66.7 lb (30.3 kg) with two drives
Warranty	Three years. IBM Ultrium media warranted separately.

12.5.5 3583 IBM Ultrium Scalable Tape Library

Figure 12-14 shows the 3583 Ultrium Scalable Tape Library.



Figure 12-14 3583 Ultrium Scalable Tape Library

The IBM TotalStorage Ultrium Scalable Tape Library 3583 is an automated tape library incorporating IBM TotalStorage LTO Ultrium Tape Drives and attaching to IBM @server iSeries, AS/400, IBM @server pSeries, RS/6000, RS/6000 SP, IBM @server xSeries, Netfinity, and other UNIX and PC servers supporting OS/400, AIX, Sun Solaris, HP-UX, Microsoft Windows NT, Microsoft Windows 2000, and Microsoft Windows 2003 open systems with a Small Computer Systems Interface (SCSI).

Model abstract 3583

This next section show the differences of the three Models:

- ▶ The IBM 3583 Model L18 has 18 cartridge slots and one to six IBM Ultrium Tape Drives.
- ▶ The IBM 3583 Model L36 has 36 cartridge slots and one to six IBM Ultrium Tape Drives.

- ▶ The IBM 3583 Model L72 has 72 cartridge slots, including a 12-cartridge I/O station, and one to six IBM Ultrium Tape Drives.

Highlights

The following are the highlights of the 3583 Ultrium Scalable Tape Library:

- ▶ Scalable from one to six IBM Ultrium drives and 18 to 72 IBM Ultrium Tape Cartridges.
- ▶ Tape drives with a SCSI Low Voltage Differential (LVD) interface, SCSI High Voltage Differential (HVD) interface, or Fibre Channel interface.
- ▶ Multi-Path Support
 - Standard for new libraries with Ultrium 2 drives
 - MES for existing libraries with Ultrium 1 drives
 - The patented Multi-Path feature of the Ultrium Tape Library 3583 supports sharing of the library robotics. Maximum of three Logical Library Partitions, providing each logical library its own separate and distinct drive(s), storage slots, and control path(s).
 - Input/output (I/O) slots are shared on a first-come-first-served basis. This type of partitioning allows heterogeneous applications to share the library robotics independent of each other.
 - Cartridges under library control are not shared between logical libraries, nor allowed to be moved between logical libraries.
 - Logical libraries can also be used for separating Ultrium 1 Tape Drives and cartridges from Ultrium 2 Tape Drives and cartridges, for applications that do not support mixing the drives in the same logical library.
- ▶ Control Path Failover for AIX Option.
 - Optional feature
 - Enables automatic failover to redundant HBA, SAN, or library control path
 - When a failure is detected by the device driver software
 - Implemented through library/drive firmware and IBM device driver updates
 - Available initially for Fibre Channel attachment
- ▶ Integration of Ultrium 2 Tape Drives with Native Fibre Channel drive option.
- ▶ Built-in bar code reader.
- ▶ External stand-alone or rack-mountable models.
- ▶ Remote Management Unit/Specialist.
- ▶ The Remote Management Unit (RMU) comes standard in every library.

- ▶ The RMU/Specialist enables network access to the library through a Web browser, for more detailed status and control.
- ▶ Library status can be sent to the network as Simple Network Management Protocol (SNMP) traps.
- ▶ All library operator panel functions can be accessed using the Specialist.

Compatibility

Here is the compatibility information:

- ▶ The IBM LTO Ultrium 2 Tape Drives can read and write original IBM LTO Ultrium Data Cartridges.
- ▶ The new LTO Ultrium 200 GB Data Cartridges can only be used on the new IBM LTO Ultrium 2 Tape Drives. However, IBM LTO Ultrium 2 Tape Drives and cartridges can be resident in the Ultrium Scalable Tape Library 3583 with IBM LTO Ultrium 1 Tape Drives and cartridges.

Configuration notes

Here are the configuration notes for the 3583 Ultrium Scalable Tape Library:

- ▶ Fibre Channel cable lengths are limited to 500 m (1650 ft.).
- ▶ Total LVD SCSI cable length is limited to 25 m (81 ft.) using point-to-point interconnection (for example, one host connected to only one tape drive). Total LVD SCSI cable length is 12 m (39 ft.) using multi-drop interconnection (for example, one host connected to more than one tape drive on the same SCSI bus). Stub length at each device must not exceed 0.1 m (0.33 ft.).
- ▶ Total HVD SCSI cable lengths are limited to 25 m (81 ft.) using point-to-point or multi-drop interconnection. Stub length at each device must not exceed 0.2 m (0.66 ft.).
- ▶ All drives installed at the plant are IBM Ultrium Tape Drives and are installed sequentially beginning at position one.
- ▶ Data transfer rates on the SCSI bus can be affected by the amount of IBM Ultrium drives attached to the bus. For optimal performance, it is recommended that no more than two IBM Ultrium drives be attached to an individual SCSI bus.
- ▶ The weight of the 3583 Scalable Tape Library is over 200 pounds. Several 3583s could be mounted in a rack enclosure and a tipping hazard could be created. Customers are advised to take necessary safety precautions when mounting multiple 3583 libraries.
- ▶ Installing more than one Ultrium drive on a SCSI bus may impact system performance. Intermixing of other SCSI devices on the same SCSI bus as the 3583 may impact performance of those devices.

- ▶ The Ultrium Scalable Tape Library 3583 operating environment must not conflict with media operating and storage requirements. If media is stored in the 3583 for more than ten hours, the media storage temperature requirements must be met.
- ▶ It is not recommended that the LTO Ultrium 2 Fibre Drive Sled (FC 8105) be installed in the same 3583 library with the San Data Gateway Module feature (FC 8005).

Table 12-15 provides summarizes the key features of the IBM TotalStorage Ultrium Scalable Tape Library 3583.

Table 12-15 IBM TotalStorage Ultrium Scalable Tape Library 3583 at a glance

Specification	Detail
Model number	<ul style="list-style-type: none"> ▶ L18 (18 cartridges) ▶ L36 (36 cartridges) ▶ L72 (72 cartridges)
Tape drive type	IBM LTO Ultrium 2 or 1
Number of Drives	Up to 6
Capacity/	Ultrium 1: Model L18: 18 Cartridges Native/Compressed: 1.8 TB/3.6 TB Model L36: 36 Cartridges Native/Compressed: 3.6 TB/7.2 TB Model L72: 72 Cartridges Native/Compressed: 7.2 TB/14.4 TB ----- Ultrium 2 Model L18: 18 Cartridges Native/Compressed: 3.6 TB/7.2 TB Model L36: 36 Cartridges Native/Compressed: 7.2 TB/14.4 TB Model L72: 72 Cartridges Native/Compressed: 14.4 TB/28.8 TB

Specification	Detail
pSeries Hardware requirements	<ul style="list-style-type: none"> ▶ FC 6204 HVD Ultra PCI ▶ FC 6207 HVD Ultra PCI ▶ FC 6205 LVD Ultra2 PCI ▶ FC 6203 LVD/SE Ultra3 (PCI) ▶ Integrated Ultra2 SCSI LVD with VHDCI ▶ Integrated Ultra3 SCSI LVD with VHDCI ▶ FC 6227 and 6228: <ul style="list-style-type: none"> – Ultrium 1 LVD Tape Drive can attach to a Fibre Channel network through an optional SAN Data Gateway Module feature (FC 8005). – Ultrium 2 Tape Drive- Switch Fabric support only. <p>Check with the sales representatives for specific host adapters supported by model and detailed attachment and configuration information.</p>

Specification	Detail
Operating systems	<p>Ultrium 1 Tape Drives in the Ultrium Scalable Tape Library 3583 are supported on the following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 4 Release 4 or later ▶ AIX Version 4.3.2 or later ▶ Sun Solaris 2.6, 7, or 8 ▶ Microsoft Windows NT 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 (build 2195 or greater) ▶ Microsoft Windows 2003 (build 3790 or greater) ▶ HP-UX 11.00 and 11i (64-bit) ▶ Linux distributions: Red Hat 7.1, 7.2, 7.3, or Advanced Server 2.1 <p>Ultrium 2 Tape Drives in the Ultrium Scalable Tape Library 3583 are supported on the following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 5 Release 1 or later ▶ AIX Version 4.3.3 or later ▶ Sun Solaris 7, 8, or 9 ▶ Microsoft Windows NT 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 (build 2195 or greater) ▶ Microsoft Windows 2003 (build 3790 or greater) ▶ HP-UX 11.00 and 11i (64-bit) ▶ Linux distributions: Red Hat Advanced Server 2.1, SuSE Linux Enterprise Server Update
Media	<ul style="list-style-type: none"> ▶ IBM LTO Ultrium Data Cartridge (100 GB) ▶ (P/N 08L9120) ▶ IBM TotalStorage LTO Ultrium 200 GB Data Cartridge (P/N 08L9870) ▶ IBM TotalStorage LTO Cleaning Cartridge (universal) (P/N 5L2086) ▶ 6-Cartridge Magazine with cover (P/N 19P4529) ▶ Leader Pin Attachment Kit (P/N 08L9129)
Dimensions	25.0 in. H x 18.9 in. W x 28.9 in. D (63.5 cm x 48.1 cm x 73.5 cm).
Weight	145 lb to 257.0 lb (65.8 kg to 116.6 kg).

Specification	Detail
Warranty	One year; IBM On-site Repair (IOR).

12.5.6 3584 IBM UltraScalable Tape Library

Figure 12-15 shows the 3584 UltraScalable Tape Library.



Figure 12-15 3584 UltraScalable Tape Library

The IBM 3584 UltraScalable Tape Library is part of a family of tape library storage solutions designed for the large, unattended storage requirements from today's mid-range systems up to high-end open systems. Each aspect of the subsystem is designed to optimize access to data and reliability. The 3584 UltraScalable Tape Library is designed to connect to host systems using any combination of Fibre Channel, Ultra2/Wide Low Voltage Differential (LVD) SCSI, or Ultra/Wide High Voltage Differential (HVD) SCSI interfaces.

Components

The 3584 is composed of a base library, Model L32 and optional expansion frames, and Model D32.

Model abstract 3584-L32

The IBM 3584 Model L32 base library has 141 to 281 cartridge slots and support for up to twelve IBM LTO Ultrium tape drives with an incremental reduction of storage slots for each set of four drives. Data capacity for the Model L32 is 14 to 56 TB native and 28 to 112 TB using LTO-DC (LTO Data Compression) (2:1) compression. Aggregate library data rates of up to 18 TB/hour (with 2:1 compression) for a fully configured 3584 library allow extremely high data transfer performance. Up to 12 logical libraries and/or up to 12 control paths can be configured for each frame. Each Model L32 library has a standard 10-slot cartridge input/output station for importing or exporting LTO cartridges from the library without requiring a re-inventory. Optional features can provide 18 input/output slots for DLTtape) IV cartridges or 20 additional input/output slots for LTO Ultrium cartridges. For bulk-loading of IBM LTO Ultrium tape cartridges, the library door can be opened. Each time the library door is closed, a bar code reader mounted on the autochanger scans the cartridge labels, enabling a re-inventory of the cartridges in the library frame in less than 60 seconds. A door lock is included to restrict physical access to cartridges in the library. Customers can expand library capacity and number of drives to meet their changing needs.

Model abstract 3584-D32

The IBM 3584 Model D32 expansion frames may be added to the base frame (Model L32) in order to add storage and/or drive capacity. Up to five D32 expansion frames may be added to each base frame. Each 3584 Model D32 frame supports up to 440 storage slots and up to twelve drives, with incremental reduction of storage slots for each set of four drives installed. Each Model D32 frame can have up to 12 logical libraries and 12 control paths.

Highlights

The following are the highlights of the 3583 UltraScalable Tape Library:

- ▶ Attachment to iSeries, AS/400, pSeries, RS/6000, xSeries, Netfinity, Sun, Hewlett-Packard, and other non-IBM servers
- ▶ Connectivity using Fibre Channel, Low Voltage Differential (LVD) SCSI, or High Voltage Differential (HVD) SCSI interfaces
- ▶ IBM multi-path architecture supporting mixed drive configurations and library sharing between multiple hosts

Model conversions

When converting the last Model D42 library to a Model D32 library within a library, RPQ 8B3246 is required for the Model L32 library. RPQ 8B3246 replaces the LTO/DLT gripper assembly in the Model L32 with a dual gripper assembly capable of handling LTO cartridges. The RPQ also removes the 18-cartridge DLT I/O station in the Model L32 and replaces it with a 20-cartridge LTO I/O station.

Configuration notes

The following are the configuration notes for the 3584 UltraScalable Tape Library:

- ▶ Although the compression technology can increase the amount of data stored on the media, the actual degree of compression achieved is highly sensitive to the characteristics of the data being compressed.
- ▶ Fibre Channel cable lengths are limited to 500 m (1640 ft.).
- ▶ Total LVD SCSI cable length is limited to 25 m (81 ft.) using point-to-point interconnection (for example, one host connected to only one tape drive). Total LVD SCSI cable length is 12 m (39 ft.) using multi-drop interconnection (for example, one host connected to more than one tape drive on the same SCSI bus). Stub length at each device must not exceed 0.1 m (0.33 ft.).
- ▶ Total HVD SCSI cable lengths are limited to 25 m (81 ft.) using point-to-point or multi-drop interconnection. Stub length at each device must not exceed 0.2 m (0.66 ft.).
- ▶ Installing more than one IBM LTO Ultrium drive on a SCSI bus may impact system performance. Intermixing of other SCSI devices on the same SCSI bus as the 3584 UltraScalable Tape Library may impact performance of those devices.
- ▶ Each IBM 3584 UltraScalable Tape Library can be divided into logical libraries (up to the number of drives).
- ▶ Each logical library must contain at least one IBM LTO Ultrium tape drive and one cartridge slot.

Table 12-16 summarizes the key features of the IBM TotalStorage UltraScalable Tape Library 3584.

Table 12-16 IBM TotalStorage UltraScalable Tape Library 3584 at a glance

Specification	Detail
Model number	L32—LTO base frame; D32—LTO expansion frame.
Tape drive type	IBM LTO Ultrium 2 or 1.
Number of frames	One base frame and up to five expansion frames.
Number of drives	Up to 72: L32—1 to 12 LTO; D32—0 to 12 LTO.
Number of tape cartridges	Up to 2481: L32—87 to 281; D32—396 to 440.
Number of logical libraries	Up to 72: L32—up to 12; D32—up to 12.

Specification	Detail
Capacity	<ul style="list-style-type: none"> ▶ 1,2 992.4 TB compressed maximum, six-frame configuration ▶ L32 (1-4 drives)—Up to 112.4 TB/frame compressed; 56.2 TB native ▶ D32 (0 drives)—Up to 176 TB/frame compressed; 88.0 TB native
pSeries hardware requirements	<ul style="list-style-type: none"> ▶ FC 6204 HVD Ultra PCI ▶ FC 6207 HVD Ultra PCI ▶ FC 6205 LVD Ultra2 PCI ▶ FC 6203 LVD/SE Ultra3 (PCI) ▶ Integrated Ultra2 SCSI LVD with ▶ VHDCI 1 ▶ Integrated Ultra3 SCSI LVD with VHDCI ▶ FC 6227 and 6228 with LTO Ultrium 2 attachment - Switch Fabric support only <p>Check with the sales representatives for specific host adapters supported by model and detailed attachment and configuration information.</p>
Operating systems	<p>Supported on the following operating systems at the minimum levels indicated:</p> <ul style="list-style-type: none"> ▶ OS/400 Version 5 Release 1 or later ▶ AIX Version 4.3.3 or later ▶ Sun Solaris 7, 8, or 9 ▶ Microsoft Windows NT 4.0 with Service Pack 6 ▶ Microsoft Windows 2000 (build 2195, or greater) ▶ HP-UX 11.00 and 11i (64-bit) ▶ Linux distributions: Linux Red Hat 7.2 32-bit and 64-bit Kernels, Linux Red Hat 7.3, Red Hat Advanced Server 2.1, SuSE Linux Enterprise Server 7 Update <p>For a current list of host software versions and release levels that support the 3584, refer to the following Web site: http://www.ibm.com/storage/lto Tivoli Storage Manager and other industry-leading compatible software offerings provides storage and tape management software for the 3584.</p>

Specification	Detail
Media	<ul style="list-style-type: none"> ▶ IBM TotalStorage LTO Ultrium 200 GB Data Cartridge (P/N 08L9870) ▶ IBM TotalStorage LTO Cleaning Cartridge (P/N 35L2086)
Dimensions	(L32 or D32) 70.9 in. H x 28.5 in. W x 59.8 in. D (1800 mm x 725 mm x 1520 mm).
Weight	Max. Weight L32—1236.0 lb (539.4 kg); D32—1161.9 lb (527.0 kg).
Warranty	Warranty One year; IBM On-site Repair (IOR).

12.6 The IBM SAN solution

The evolution of information technology (IT), coupled with the explosion of both the intranet and Internet, has led to a storage capacity need that doubles every quarter. This goes on year by year and leads to the need for flexible and fast access to the servers, and the data stored on fast, accessible storage. To achieve this, we need a network that is fast, multipathing, easy to manage, and available 24x7, 365 days a year. The answer to this is a SAN. A SAN will enable host servers to make use of any storage device (disk or tape) attached to the SAN. SANs promise an expansion of openness by allowing heterogeneous storage and server attachment to the SAN. In theory, the user will have the freedom to attach storage to the SAN independent of server type or storage vendor. SAN management, using an abstraction of storage that is called virtualization, will further break the traditional server/storage ownership relationship. For more information on SAN for pSeries, see the *Practical Guide for SAN with pSeries*, SG24-6050. The elements of a SAN can be summarized into the following:

- ▶ Server operating systems and host bus adapters
- ▶ Fibre Channel fabric interconnect components
- ▶ Storage subsystems and/or devices
- ▶ Storage abstraction by appliance or software solution
- ▶ SAN management software
- ▶ Software to exploit SAN capabilities

12.6.1 IBM @server pSeries Fibre Channel Host Bus Adapters (HBA)

The following section describes the Fibre Channel host bus adapters (HBA) used for pSeries.

FC 6228 2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus

(For IBM US, No Longer Available as of August 8, 2003)

The 2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus is a 64-bit address/data, short form factor PCI adapter with LC type external fiber connectors that provides single or dual initiator capability over an optical fiber link or loop. With the use of appropriate optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage. The 2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus will auto-negotiate for the highest data rate (either 1Gps or 2Gps) of which the device or switch is capable. Distances of up to 500 meters running at 1 Gbps data rate and up to 300 meters running at 2 Gbps data rate are supported between the adapter and an attaching device or switch. When used with IBM supported Fibre Channel storage switches supporting long-wave optics, distances of up to 10 kilometers are capable running at either 1 Gbps or 2 Gbps data rates.

The 2 Gigabit Fibre Channel Adapter for 64-bit PCI Bus can be used to attach devices either directly, or by means of Fibre Channel Switches. If attaching a device or switch with a SC type fiber connector(s), use of an LC-SC Fibre Channel Conversion Cable (FC 2456) is required.

Refer to the following IBM storage subsystem Web page for additional supported server attachment information for IBM devices:

<http://www.storage.ibm.com/hardsoft/disk/products.htm>

FC 6239 - 2 Gigabit Fibre Channel PCI-X Adapter

The 2 Gigabit Fibre Channel PCI-X Adapter is a 64-bit address/data, short form factor PCI-X adapter with LC type external fiber connectors that provides single or dual initiator capability over an optical fiber link or loop. With the use of appropriate optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage. The 2 Gigabit Fibre Channel PCI-X Adapter will auto-negotiate for the highest data rate (either 1 Gbps or 2 Gbps) of which the device or switch is capable. Distances of up to 500 meters running at 1 Gbps data rate and up to 300 meters running at 2 Gbps data rate are supported between the adapter and an attaching device or switch. When used with IBM supported Fibre Channel storage switches supporting long-wave optics, distances of up to 10 kilometers are capable running at either 1 Gbps or 2 Gbps data rates.

The 2 Gigabit Fibre Channel PCI-X Adapter can be used to attach devices either directly, or by means of Fibre Channel Switches. If attaching a device or switch with a SC type fiber connector(s), use of an LC-SC 50 Micron Fiber Converter Cable (FC 2456) or a LC-SC 62.5 Micron Fiber Converter Cable (FC 2459) is required.

Refer to the following IBM storage subsystem Web page for additional supported server attachment information for IBM devices:

<http://www.storage.ibm.com/hardsoft/disk/products.htm>

12.6.2 IBM TotalStorage SAN switches

The various models of the IBM TotalStorage SAN Fibre Channel Switch 2109 and 3534-F08 provide Fibre Channel connectivity to a large variety of Fibre Channel-attached servers and disk storage, including the IBM TotalStorage Enterprise Storage Server (ESS), FASTT Storage Servers, SAN Data Gateways for attachment of IBM Enterprise Tape systems 358x, and tape subsystems with native Fibre Channel connections. This chapter give some details about these products and describe their interactions.

IBM offers four different models of switches, which are OEM products from the Brocade SilkWorm family, as follows:

- ▶ The IBM TotalStorage SAN Fibre Channel Switch Model 3534-F08 is an 8-port model.
- ▶ The IBM TotalStorage SAN Fibre Channel Switch Model 2109-F16 is a 16-port model.
- ▶ The IBM TotalStorage SAN Fibre Channel Switch Model 2109-F32 is a 32-port model.
- ▶ The IBM TotalStorage SAN Fibre Channel Switch Model 2109-M12 is a (dual) 64-port model.

3534-F08

While the older IBM SAN Switches 2109-S08 and 2109-S16 supported ANSI standard Fibre Channel protocol at 1 Gbps, all the new models are built upon a third-generation switch technology that supports a link bandwidth of 1 and 2 Gbps. These third-generation or next-generation switches are often referred to as “2-Gbps switches”.

The ports of all of the IBM SAN switches are numbered sequentially, starting with zero for the left-most port. The switch faceplate includes a silk screen imprint of the port numbers. With the 2 Gbps switches, the ports are color-coded into quad-groups to indicate which ports can be used in the same ISL trunking group.

The 3534-F08, as shown in Figure 12-16, is an 8-port SAN switch. It supports Fibre Channel classes 2, 3, and F, and has a latency of less than 2 μ s with no contention (assuming the destination port is free).



Figure 12-16 IBM SAN Fibre Channel Switch 3534-F08

Classes of service: Class F is a connectionless service for inter-switch control traffic. It provides notification of delivery or nondelivery between two E_Ports. Class 2 is a connectionless service between ports with notification of delivery or non-delivery. Class 3 is a connectionless service between ports without notification of delivery. Other than notification, the transmission and routing of Class 3 frames is the same as Class 2 frames.

The basic configuration of the 3534-F08 provides four SW-SFPs. In addition, a mixture of SW and LW ports can be configured by adding up to eight SFP transceivers. High availability is supported by hot-pluggable cooling fans and SFPs. An additional power supply can be ordered for redundancy.

The Entry Fabric switch configuration includes a fabric software license and connection to one other F08, and comes without the zoning feature. The Full Fabric upgrade option is needed to implement zoning. Web Tools is also included in the basic configuration.

Additional software features available are:

- ▶ Full Fabric Activation (FC 7320)
- ▶ Performance Bundle (C 7321)
- ▶ Extended Fabric (FC 7303)
- ▶ Remote Switch (FC 7302)

For more information, 3534-F08 data sheets can be downloaded from:

<http://www.storage.ibm.com/ibmsan/products/2109/library.html#support>

2109-F16 and 2109-F32

The F16 and F32 are also built upon the third-generation switch technology that supports link bandwidth of 1 and 2 Gbps. The F16 is shown in Figure 12-17 and is a 16-port switch.




Figure 12-17 IBM SAN Fibre Channel Switch 2109-F16

The ports of all of the IBM SAN switches are numbered sequentially, starting with zero for the left-most port. The switch faceplate includes a silk screen imprint of the port numbers. With the 2 Gbps switches, the ports are color-coded into quad-groups to indicate which ports can be used in the same ISL trunking group.

The F16 consists of a system board with connectors for supporting up to 16 ports and a Fabric Operating System for building and managing a SAN fabric. The F16 supports Fibre Channel classes 2, 3, and F, and has a latency of less than 2 μ s with no contention (assuming the destination port is free).

The base model F16 configuration comes with eight SW or LW-SFPs. In addition, a mixture of SW and LW ports can be configured by adding up to sixteen SFP transceivers. High availability is supported by hot-pluggable cooling fans and SFPs, and an additional power supply can be ordered for redundancy.

Advanced Zoning and Web Tools are included in the F16 basic configuration.

Additional software features available are:

- ▶ Performance Bundle (FC 7421)
- ▶ Extended Fabric (FC 7303)
- ▶ Remote Switch (FC 7302)
- ▶ Fabric Manager (FC 7202)

The F32, as shown in Figure 12-18 on page 516, is a 32-port switch that shares the same characteristics as the F16.



Figure 12-18 IBM SAN Fibre Channel Switch 2109-F32

The base model F32 configuration comes with 16 SW-SFPs. In addition, a mixture of SW and LW ports can be configured by adding up to sixteen SFP transceivers. High availability is supported by hot-pluggable cooling fans and SFPs. A redundant power supply enabling dual-power and non-disruptive power supply maintenance is included.

The Performance Bundle function, Advanced Zoning, and Web Tools are also included in the F32 basic configuration.

Additional software features available are:

- ▶ Extended Fabric (FC 7303 for F16, FC 7503 for F32)
- ▶ Remote Switch (FC 7302 for F16, FC 7502 for F32)
- ▶ Fabric Manager (FC 7202)

For more information, the F16 and F32 data sheets can be downloaded from:

<http://www.storage.ibm.com/ibmsan/products/2109/library.html#support>

2109-M12

The M12 is a bladed architecture, which consists of two logical 64-port switches in one chassis. Each logical switch has its own:

- ▶ Unique domain ID
- ▶ Switch World Wide Name (WWN)
- ▶ IP address

Both logical switches share:

- ▶ Chassis
- ▶ Four hot swappable, redundant power supplies:
 - Any two are needed to provide power for a maximum configuration.

- Selective power down of cards (configurable) when only one power supply is working.
- Two AC connectors.
- ▶ Three hot swappable, redundant fans
 - Any two needed to cool the entire switch.
 - One fan can keep the unit running for about one hour.
- ▶ Dual control processor (CP) operates in active/standby mode

The M12 is shown in Figure 12-19.



Figure 12-19 2109-M12

12.6.3 Sample configurations

This section shows some configurations with pSeries, switches, and external storage devices.

Small configuration

Figure 12-20 on page 518 show a schematic small SAN solution with:

- ▶ Two IBM @server pSeries 615s, each with one FC adapter (FC 6239)
- ▶ One 3534 F08 switch
- ▶ One 3553 tape

- ▶ One FAStT600

This configuration is not recommended for high availability.

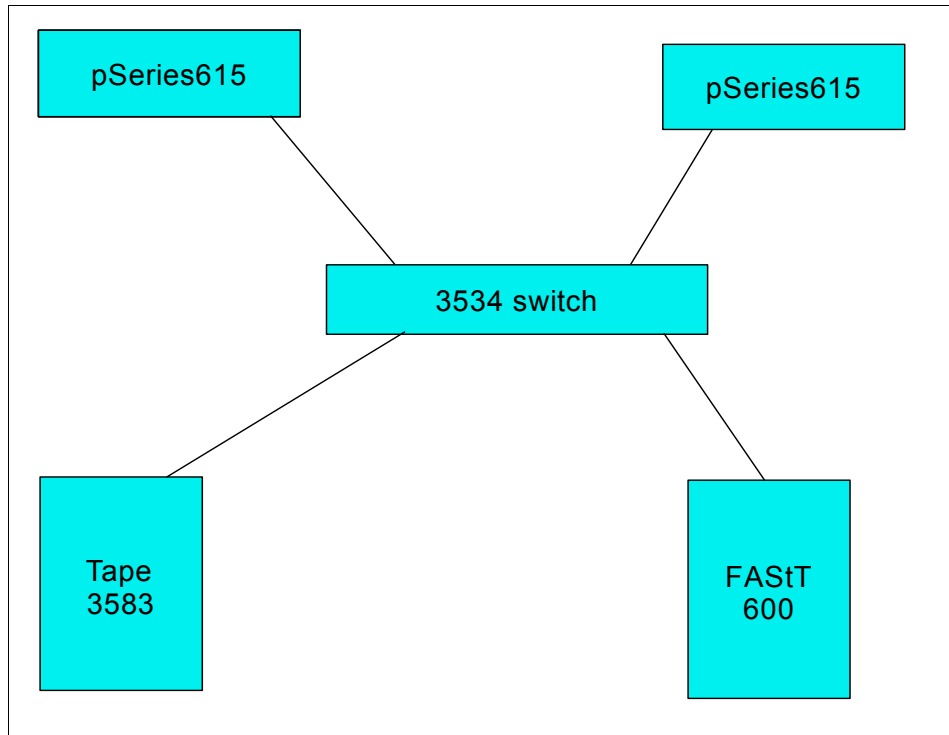


Figure 12-20 Schematic small SAN solution

Medium configuration

Figure 12-21 on page 519 shows a schematic medium SAN solution with:

- ▶ Two IBM @server pSeries 650s, one with LPAR, each with two FC adapters (FC 6239)
- ▶ Two 2109 F16 switches
- ▶ One 3584 tape
- ▶ One FAStT900

This is a high availability configuration together with HACMP software and at least two network adapters on each system or LPAR.




Figure 12-21 Schematic medium SAN solution

Large configuration

Figure 12-22 on page 520 show a schematic large SAN solution with:

- ▶ Two IBM @server pSeries 690s, one with LPAR, each with minimum of two FC adapters (FC 6239)
- ▶ Two 2109-F32 switches
- ▶ Two ESS 2105 800s
- ▶ Two 3494 tape libraries

This is a high availability configuration together with HACMP software and at least two network adapters on each system or LPAR.




Figure 12-22 Schematic large SAN solution



Operating systems

This chapter discusses the AIX 5L Version 5.1 and Version 5.2 and Linux operating systems, which provide the interface to interact with the pSeries systems for the user and applications.

13.1 AIX 5L Version 5.1

Advanced Interactive Executive (AIX) 5L represents the next generation of AIX and is an open, enterprise class UNIX operating system incorporating technology from software developers around the world. Built on a proven code base, AIX 5L is designed to exploit advanced 64-bit system and software architectures while introducing:

- ▶ Logical partitioning
- ▶ Improved workload management
- ▶ Integrated Linux affinity
- ▶ Network performance improvement
- ▶ System security enhancements
- ▶ Reliability, availability, and serviceability (RAS) enhancements and performance-tuning tools
- ▶ Cluster Systems Management

This section introduces general characteristics and features of AIX 5L Version 5.1.

13.1.1 Hardware requirements

AIX 5L Version 5.1 operates on IBM Power, POWER2, Personal Computer Power Series 830 and 850 desktop systems, IBM PowerPC systems, POWER3, POWER4, or POWER4+ systems with the following exceptions:

- ▶ RS/6000 7016 POWERserver® Model 730
- ▶ RS/6000 7007 Notebook Workstation Model N40
- ▶ POWERnetwork Dataserver 7051
- ▶ RS/6000 7249 Models 851 and 860
- ▶ RS/6000 7247 Models 821, 822, and 823

AIX 5L Version 5.1 supports system with at least 64 MB of physical memory, 128 MB of initial disk paging space, and requires 536 MB disk storage for the operating system for a total of 664 MB of disk storage.

All POWER graphics adapters supported on AIX Version 4.3.3 will also be supported on the AIX 5L Version 5.1 32-bit kernel.

Additionally, the following POWER graphics adapters will be supported on both 32-bit and 64-bit kernels:

- ▶ POWER GXT120P
- ▶ POWER GXT130P
- ▶ POWER GXT300P
- ▶ POWER GXT2000P
- ▶ POWER GXT3000P
- ▶ POWER GXT4000P
- ▶ POWER GXT6000P

OpenGL on POWER GXT4000P and GXT6000P graphics adapters will now support 64-bit direct window access (DWA). This is intended to boost performance for 64-bit OpenGL applications by allowing them to render using the OpenGL protocol directly, rather than going through the xServer and GLX Extension.

AIX 5L Version 5.1 will *not* support the following hardware features:

- ▶ 3Com 10/100 Mbps PCI Fast Etherlink XL for Power PC Systems (FC 2986)
- ▶ Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (FC 2708)

AIX 5L Version 5.1 does *not* support the following graphic input devices with a 64-bit kernel; the devices are supported under 32-bit kernel:

- ▶ 6094-010 Dials
- ▶ 6094-020 LPFK
- ▶ 6093-011 Tablet
- ▶ 6093-012 Tablet
- ▶ 6093-021 Tablet

13.1.2 Software requirements

Table 13-1 outlines the memory and disk requirements of the various software. The sections that follow discuss these requirements in detail.

Table 13-1 Memory and disk requirements

Software	Memory requirement	Disk requirement
AIX 5L Version 5.1	Minimum 64 MB	128 MB of initial disk paging space + 536 MB disk storage for the operating system = Total of 664 MB of disk storage
OPENGL and GL 3.2	<ul style="list-style-type: none">▶ 32-bit kernel: Minimum of 64 MB▶ 64-bit kernel: Minimum of 64 MB	5 MB - 50 MB
PHIGS	<ul style="list-style-type: none">▶ 32-bit kernel: Minimum of 64 MB▶ 64-bit kernel: Minimum of 64 MB	10 MB - 120 MB

SecureWay® Directory Server

To install the IBM SecureWay Directory Server, the server requires the following:

- ▶ A minimum of 64 MB RAM (128 MB is strongly recommended)
- ▶ One of the following Web servers (or a later version), installed and configured:
 - IBM HTTP Server 1.3.12
 - Lotus® Domino™ Go Web server 4.6.2.6
 - Lotus Domino Enterprise 5.0.2b Web server
 - Apache Server 1.3.12
 - Netscape FastTrack Server 3.01
 - Netscape Enterprise Server 3.6.3, 4.0

Attention: You must have a secure Web server to ensure that the transmission of your administration data is secure. For instructions, consult the online documentation.

- ▶ DB2® Version 6.1 is included with the IBM SecureWay Directory. Higher versions of DB2 might also be supported. If you already have DB2 installed, you need approximately 45 MB of disk space. You need approximately 135 MB of disk space for both LDAP and DB2. Disk space required for data storage is dependent upon the number and size of database entries.

13.1.3 AIX 5L Version 5.1 release and support date

Table 13-2 describes information of release and withdrawal date of AIX 5L Version 5.1

Table 13-2 Release and withdrawal date

Program Number	AIX Release level	Available date	Customer service withdrawal date
5765-E61	AIX 5L Version 5.1	May 04, 2001	April 01, 2006

13.1.4 AIX 5L Version 5.1 highlights

AIX 5L Version 5.1 offers:

- ▶ Full 64-bit kernel, device drivers, and application environment
- ▶ Support for 32-Way scalability and 256 GB memory
- ▶ New JFS2 1 TB file system with 1 TB file size support
- ▶ Workload Manager GUI and functional upgrades
- ▶ Increased network performance for e-business
- ▶ Reliable Scalable Cluster Technology
- ▶ Enhanced RAS and improved serviceability features
- ▶ Web-based System Manager distributed framework
- ▶ Linux-compatible program interface
- ▶ AIX Developer Kit, Java Technology Edition, Version 1.3.0

13.1.5 Scalability and system management

AIX 5L provides many enhancements in the area of scalability and system management and utilities. This section discusses these enhancements.

64-bit kernel

In addition to providing a 32-bit kernel, AIX 5L Version 5.1 offers a scalable, 64-bit kernel capable of supporting increased system resources and much larger application workloads on 64-bit hardware. The 64-bit kernel offers scalable

kernel extension interfaces, allowing kernel extensions and device drivers to make full use of the kernel's system resources and capabilities.

The expanded capabilities of the 64-bit kernel improve the ability to run an expanding application workload on a single system. This ability is important for a number of reasons.

- ▶ First, data sharing and I/O device sharing are simplified if multiple applications can be run on the same system.
- ▶ Second, using more powerful systems will reduce the number of systems needed by an organization, reducing the cost and complexity of system administration.

Server consolidation and workload scalability will continue to require higher capacity hardware systems that support more memory and additional I/O devices. The 64-bit kernel is designed to support these requirements. Kernel extensions and device drivers must be compiled in 64-bit mode to be loaded into the 64-bit kernel. The 64-bit kernel, combined with header files and libraries, provides the environment for porting and developing kernel extensions.

64-bit application scalability

AIX 5L Version 5.1 provides a more scalable application binary interface (ABI) for 64-bit applications. This scalability is provided by changing the sizes of some fundamental data types for 64-bit applications, and will allow these applications to take advantage of the expanded capabilities of the 64-bit kernel. The scalable 64-bit ABI is supported by the 32-bit kernel and the 64-bit kernel.

Web-based System Manager

Web-based System Manager for AIX 5L Version 5.1 represents a significant advance over previous releases by providing a new management console capable of managing multiple hosts. As in previous releases, Web-based System Manager applications can be accessed from PC clients running Web-browsers, such as Netscape Navigator or Microsoft Internet Explorer, that accept the Java plug-in. These applications can also be accessed either locally or remotely without a Web-browser on AIX 5L Version 5.1 graphical workstations. In addition, Web-based System Manager delivers the capability to run Window PC clients natively, significantly increasing the Web-based System Manager's performance.

New functions include:

- ▶ RSCT is the Reliable Scalable Cluster Technology offering that provides the capability for an administrator to monitor resources on the machine, including file systems, programs, processors, adapters, and kernel information, such as CPU statistics and memory allocation.

These resources can be monitored around the clock and automatic responses run if the values of resources reach certain levels or thresholds. The automated responses include running any command or script (for example a user-provided pager command or a recovery script), broadcasting a message, sending e-mail, or logging the event to a file. Different responses can be run based on time of day and whether a particular problem occurred or has been resolved.

There is a command line interface for this functionality and a graphical user interface with the Web-based System Manager monitoring application.

- ▶ A new unified management console for system administration: Web-based System Manager provides a single console and application suite.
- ▶ Enhanced scalability of the graphical user interface: Web-based System Manager includes a number of features to deal with large numbers of managed objects.
- ▶ Enhanced ease of management and usability: Web-based System Manager further simplifies administrative tasks through improvements in task design, new user interface features, and enhanced user assistance technology. This version also improves the accessibility of the console to users with disabilities.
- ▶ Simplified log on: The console and management infrastructure provides features to reduce the need to log on to multiple systems.
- ▶ Persistent and more flexible user customization options: Administrator preferences such as the choice of managed resources that are presented, can be customized and are persistent over management sessions.

Web-based System Manager application support

A set of Web-based System Manager plug-ins are available. The Web-based System Manager provides applications for all major AIX system management tasks. All of the application functions of prior releases is included with new functions and enhancements. The complete set of application functions will include:

- ▶ Host overview
- ▶ Backups
- ▶ Custom tools
- ▶ Devices
- ▶ File systems
- ▶ Monitoring
- ▶ Network
- ▶ Network Install Manager (NIM)

- ▶ PC services
- ▶ Printers
- ▶ Processes
- ▶ Software
- ▶ System environments
- ▶ Users and groups
- ▶ Volumes
- ▶ Workload Manager

Figure 13-1 shows one of WSM features that system administrators can use to perform tasks easily. In this example, WSM provides easy-access interface to administrators for seeing and modifying tuning parameters of virtual memory.

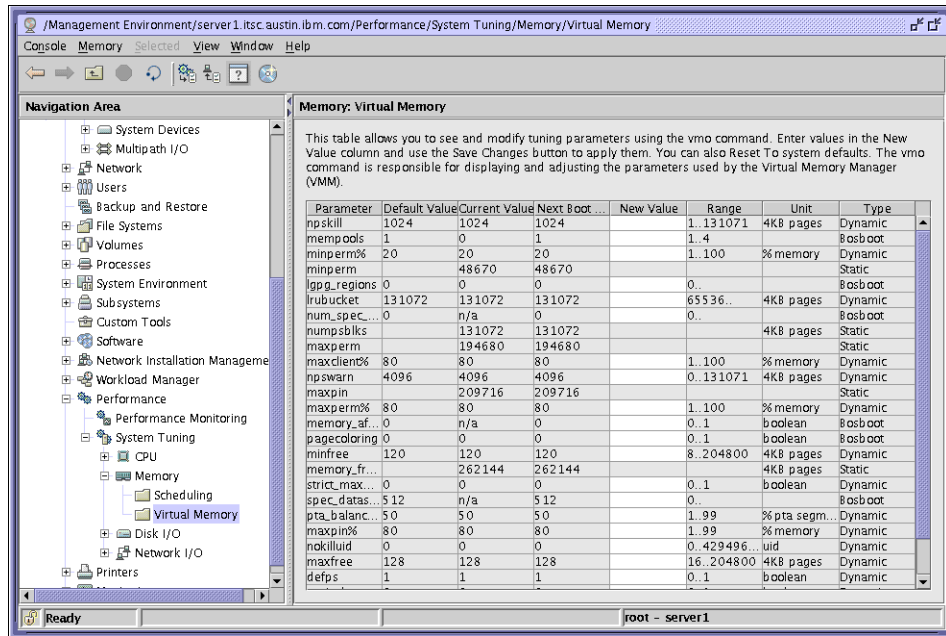


Figure 13-1 Example of Web-based System Management

Accessibility enhancement to WSM framework

Accessibility enhancement includes the following features:

- ▶ Full keyboard accessibility support, including keyboard navigation, menu shortcuts keys, and mnemonics in menu and dialogs, allowing the user to use Web-based System Manager without a mouse

- ▶ Available text descriptions for all icons and labels in Web-based System Manager
- ▶ Enablement in Web-based System Manager for speech output, using IBM Self-Voicing Kit (SVK)

This toolkit can be downloaded from the IBM AlphaWorks site:

<http://www.alphaworks.ibm.com/tech/svk>

These accessibility features can increase ease of use for all users. For example, keyboard shortcuts speed interaction for power users.

International language support

AIX offers a single worldwide desktop system offering an internationalized and localized environment that enables development and usage of global applications for worldwide markets. The internationalization (NLS) support extends across all components of the system, including basic and graphical operating system, graphical user interface, and communications.

Major NLS categories provided are for character handling, character encoding, support for language-dependent functions, character usage in object names, and language-dependent prompts, screens, panels, and messages, in accordance with national standards and cultural conventions by language territory (locale).

To assist in classifying the extent to which a separately purchasable AIX licensed program provides National Language Support, the following levels of National Language Support Enablement can be found in the program product announcement material:

- ▶ Full International Language Support (ILS)
Support for *all* language territories (locales) in the underlying operating system.
- ▶ Multi-Byte Character Set Support (MBCS)
Support for language territories (locales) is based on multi-byte and single byte code sets in the underlying operating system. Bidirectional code set support is limited.
- ▶ Single Byte Character Set Support (SBCS)
Support is limited to single byte language territories (locales) in the underlying operating system. Bidirectional code set support is limited.
- ▶ PASSTHRU Support
PASSTHRU Support is limited to the ability of the LP to flow data through the program without processing the information in such a manner that all data, control, and graphic characters flow unaltered through the program to another program

AIX provides a wide range of language territory (locale) support, as shown in Table 13-3.

Table 13-3 *Locale support*

Code	Language
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. <ul style="list-style-type: none"> ▶ 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)
Bidirectional (BIDI)	Arabic, Hebrew
Multi-Byte Character Sets (MBCS)	Traditional Chinese, Japanese (Kanji) Korean, Simplified Chinese
UTF-8	Estonian, Lithuanian, Latvian, Simplified Chinese

The basis for international language support is a product that 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.

Code Page 943

Code Page 943 supports interoperability with Microsoft Windows clients in Japan.

Euro currency symbol support

AIX strategic euro currency symbol support is based on the new Unicode locales. The existing national currency symbols and the euro currency symbol are available. Characters will appear in the proper order in the locale's collating sequence (as defined by National Language Design Guide (NLDG) Volume 2, and as amended by National Language Technical Center (NLTC) in Toronto). Applications that adhere to standard programming interfaces and are internationalized will be euro-enabled once the new euro-enabled localization is installed. Unicode locales that will have euro currency symbol support are described in Table 13-4.

Table 13-4 Euro currency symbol

Language	Locale
Catalan	CA_ES
Dutch (Belgium)	NL_BE
Dutch	NL_NL
Finnish	FI_FI
French (Belgium)	FR_BE
French	FR_FR
German	DE_DE
Italian	IT_IT
Portuguese	PT_PT
Spanish	ES_ES

These locales will include keyboard definitions using the euro currency symbol (based on the European Commission Recommendation). The keyboard support is limited to those countries within the European Monetary Union. AIX will provide code conversion capabilities for the new code sets defined to include the euro currency symbol.

Internationalized Classes for Unicode

AIX 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.

ISO Standard ISO8859-15

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. Supported locales include Italian-Switzerland (it_CH or IT_CH), English-Australia (en_AU or EN_AU), English-Belgium (en_BE or EN_BE), and English-South Africa (en_ZA or EN_ZA).

New locales in AIX 5L Version 5.1

New locales are introduced in AIX 5L, as shown in Table 13-5. Most of these are variations of existing locales, but may have slight differences from a neighbor.

Table 13-5 New locales in AIX 5L Version 5.1

Languages	Locales
Arabic	ar_AE, ar_BH, ar_EG, ar_JO, ar_KW, ar_LB, ar_OM, ar_QA, ar_SA, ar_SY, and ar_TN
English	en_CA, en_IN, en_IE, en_IE EURO, and en_NZ
French	fr_LU, and fr_LU EURO
German	de_AT, de_AT EURO, de_LU, and de_LU EURO
Serbian	sr_YU and sh_YU
Spanish	es_AR, es_CL, es_CO, es_MX, es_PE, es_US, es_UY, and es_VE
Hindi	HI_IN

Enhancements to input method editor for GBK locale

AIX 5L Version 5.1 supports additional popular input method editors such as Intelligent ABC, Pinyin, BiaoXing Ma, Internal Code, and ZhengMa for GBK codes on AIX GBK locale.

NFS statd multithreading

The status monitor provides a general framework for collecting network status information. Implemented as a daemon that runs on all NFS configured machines, the status monitor provides a simple protocol that allows applications to easily monitor the status of other machines.

Deactivating active paging spaces

This function provides new flexibility without rebooting after:

- ▶ Changing configurations
- ▶ Moving paging space to another drive
- ▶ Dividing paging space between drives

Until this release, allocated and activated paging space had to remain active until the next reboot. With this release, paging space can be deactivated without rebooting by using the new **swapoff** command. A new **shrinkps** command:

- ▶ Creates a new, temporary space
- ▶ Deactivates the original
- ▶ Changes the original to be smaller and reactivates it
- ▶ Deactivates the temporary space and returns it to logical volume status

The use of a shell script reduces the possibility of an unbootable state because users will not be allowed to run out of adequate paging space. The script checks paging space actually in use and adds a buffer for a paging space warning threshold.

Malloc enhancements

AIX 5L Version 5.1 provides an optional buckets-based extension of the default memory allocator (the malloc subsystem) that improves performance for applications that issue large numbers of small allocation requests. Each bucket consists of a block of memory that is subdivided into a predetermined number of smaller allocatable blocks of uniform size. Organizing allocatable memory in this fashion often provides faster access for allocation requests falling within the range of sizes defined by the buckets. When this capability is enabled, allocation requests that fall within a predefined range of block sizes are processed by buckets. All other requests (for example, those outside the defined range of sizes) are processed in the usual manner by the default allocator. Up to 128

buckets are available per heap. Number of buckets, bucket sizing factor, and other configuration values are specified using an environment variable prior to process startup. More information on configuring the MALLOCCTYPE and MALLOCBUCKETS environment variables is available in the book *AIX 5L Version 5.1 General Programming Concepts: Writing and Debugging Programs*, found at:

http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprgdd/genprog/genprogctfrm.htm

SVR4 printing subsystem

AIX 5L Version 5.1 provides the UNIX System V style file spooling subsystem and makes it available as an administrator configurable option. Enabling this option allows users who are more comfortable with System V printer utilities to more easily use AIX 5L Version 5.1.

Multithreaded autoFS

AutoFS is enhanced to support the multi-threaded automountd daemon. This allows automountd to provide more jobs than before. In addition to fulfilling mount requests, the new automountd handles requests such as:

- ▶ Lookup
- ▶ Unmount requests
- ▶ Autofs readdir

It is 64-bit enabled, which permits it to be run in either 32- or 64-bit kernel.

The /proc file system

The /proc file system provides access to the state of each process and thread in the system. The file system is mounted over /proc. Standard system call interfaces, such as open(), read(), write(), lseek(), are used to access /proc files. Programs such as debuggers can use /proc to control a process being debugged. /proc provides the ability to:

- ▶ Stop and start threads in a process
- ▶ Trace syscalls and signals
- ▶ Read and write virtual memory in a process
- ▶ Other capabilities

Increased concurrent group per process limit

The number of concurrent groups per process has been increased from 32 to 64. The concurrent group set is used to control access to files and programs. Increasing this value allows users greater flexibility in classifying and protecting

system resources. Administrators who require access to resources belonging to a large number of groups will be able to access those resources without having to change their concurrent group set.

Increased argument list limit

The argument list limit has been increased to 512 KB from 24 KB (24 KB is the traditional AIX limit and system default). Applications can specify the length of command line argument and the length of the environment (stream length) list in bytes. Users will be able to use the default or automatically specify an upper limit (512 KB) for their application.

Very large program support

AIX 5L Version 5.1 offers new flexibility for 32-bit maxdata programs by allowing the segments of the data heap (up to eight 256 MB segments) to be created dynamically. The large program support previously available did not allow data heap segments to be used for any other purpose, even if the data heap never grew large enough to use all the reserved segments. The very large program support will allow the data heap segments to be created dynamically. Until a segment is needed for the data heap, it may be used by shmat and mmap. In addition, when very large program support is enabled for an application, segments allocated by shmat and mmap are allocated in descending order instead of ascending order.

ANSI terminal support

The terminal emulation in AIX 5L has been enhanced to support the ANSI terminal type. The telnet session negotiates a terminal type of ANSI, so the TERM environment variable gets set to TERM=ansi. This helps reduce problems when opening a SMIT screen.

13.1.6 Resource management

In this section, we will discuss resource management of AIX 5L.

Workload Manager enhancement

Workload Manager has been enhanced to provide functions to manage subsets of workload and control subsets of total system resources. AIX Workload Manager includes the following capabilities, which can be easily managed through Web-based System Manager, SMIT, shell scripts, or the command line interface:

- ▶ Provides the continuation of those functions required to manage subsets of workload.
- ▶ Controls subsets of total system resources.

- ▶ Adds more sophistication to the externals for categorization of work in the system and for the specification of policy.
- ▶ Uses an alternative approach to divide up system resources and schedule a portion of the installation's total workload against a subset of the system resources.
- ▶ Disk I/O Bandwidth, a new resource type, is introduced, in addition to existing resources, such as CPU cycles and real memory.
- ▶ An Application Programming Interface (API) enables external applications to modify system behaviors.
- ▶ System administrators can manually reclassify processes independent of the classification rules. This function enables multiple instances of the same application to exist in different classes. Using Application Tag API, applications can enable automatic assignment of multiple instances of the same application in different classes.
- ▶ More application isolation and control are offered:
 - New subclasses add ten times the granularity of control (from 27 classes to 270 controllable classes).
 - System administrators can delegate subclass management to users or groups.
- ▶ Fully dynamic, Workload Manager allows an entire configuration to be changed while it is running.
- ▶ Application path name extends wild card flexibility to user name and group name.
- ▶ The accounting subsystem, a new feature, allows users to perform resource usage accounting per WLM class in addition to the standard accounting per user or group.
- ▶ The AIX accounting system utility allows system administrators to collect and report the use of various system resources by user, group, or WLM class.
 - When process accounting is turned-on, AIX records statistics about the process resource usage in an accounting file when the process exits. (This accounting record now includes a 64-bit numeric key representing the name of the WLM class and the process belonged to).
 - The accounting command, **acctcom**, allows the display of process resource usage statistics per user, group, or WLM class.
- ▶ System administrators can use either **acctcom** or the WLM API to write their own reporting and billing applications.

System administrators also can have the CLASS name associated with CPU, memory, and disk I/O consumption for any given user-defined time period.

13.1.7 Storage management

This section describes storage management of AIX 5L.

DMAPI (Data Management API)

AIX 5L Version 5.1 includes the implementation of the Data Management Application Programming Interface (DMAPI), as defined in the Open Group's Data Storage Management (XDSM) API specification. DMAPI is a technology that enables the development of relatively portable data management applications, such as hierarchical storage management (HSM), by providing the underlying JFS support and programming interface for those applications. DMAPI is available on the 32-bit kernel only.

LVM scalability (Variable LTG)

AIX 5L now supports different logical track group (LTG) sizes. In previous versions of AIX, the only supported LTG size was 128 KB. This is still the default for the creation of new volume groups, even under AIX 5L. You can change this value when you create a new volume group with the **mkvg** command, or later for an existing volume group with the **chvg** command.

The LTG corresponds to the maximum allowed transfer size for disk I/O (many disks today support sizes larger than 128 KB). To take advantage of these larger transfer sizes and get better disk I/O performance, AIX 5L now accepts values of 128 KB, 256 KB, 512 KB, and 1024 KB for the LTG size, and possibly even larger values in the future.

LVM mirror write consistency

Mirror Write Consistency (MWC) helps ensure data integrity on logical volumes in the event of a system crash during mirrored writes. The existing method achieves this by logging when a write operation occurs. The LVM makes an update to a log (MWC log) that identifies what areas of the disk are being updated prior to performing the write of the data. This results in a performance degradation during random writes. There are now two ways of handling MWC, which are active and passive method. The existing method is Active MWC and the new MWC is called passive. The *Active MWC* method (existing MWC) is still supported and is the default mode when a mirrored logical volume is created. *Passive MWC* reduces the instance of having to update the MWC log on the disk. This method logs that the logical volume has been opened but does not log the writes. If the volume group is not deactivated before reboot (crash), then the LVM starts a forced synchronization of the entire logical volume when the system restarts. Data consistency exists for reads that occur during the synchronization so that applications can start using data as soon as the volume group is varied on. The new Passive MWC method is intended to provide better random write

performance on mirrored logical volumes when compared to the Active MWC method.

Hot spare disk support in volume group

The Logical Volume Manager (LVM) previously shipped with AIX took no action when partitions went stale or disks went missing other than logging an error in the error log. The new ability of LVM to optionally provide hot spare disks within a volume group and automatically synchronize stale partitions helps to increase the availability of mirrored data. In addition, the Logical Volume Manager will attempt to reactivate a missing disk. If the disk cannot be restarted, then its mirror copy will be migrated to a hot spare disk if such a disk with the proper size exists.

Hot spot management in logical volume group

LVM provides the ability to move physical partitions (smallest possible data unit in a volume group) to any member disk of the volume group. Sometimes it is necessary to migrate partitions to a new disk for maintenance (for example, disk replacement). This migration may also achieve a performance gain for I/O if there happened to be high-traffic partitions all located on a single disk and some were migrated to other member disks. Currently, there are no tools at LVM level that will identify the partitions by the number of I/O structures. There are tools that provide similar information like `iostat` and `filemon`, but they do not directly identify the partitions at the LVM level. Hot spot management provides two commands: one that will identify the hot spots and another that migrates the hot spot to a different location.

Very large file support

Cachefs is enhanced to work under the 64-bit kernel and will support large file systems. It will handle files larger than 2 GB, although the cache file does not have to be larger than 2 GB. As a result, data can be accessed faster.

JFS2 support

JFS2 is a new file system providing the capability to store much larger files than the existing Journaled File System (JFS). It is the default file system for the 64-bit kernel. Customers have the choice of implementing JFS or taking advantage of the additional 64-bit functionality of JFS2. Table 13-6 on page 539 provides a summary of the differences between JFS2 and JFS.

Table 13-6 Comparison between JFS2 and JFS

Functions	JFS2	JFS
Fragments/Block size	512-4096 Block size	512-4096 Fragments
Architectural Max. file size	4 PB	64 GB
Maximum file size tested	1 TB	64 GB
Architectural Max. file system size	4 PB	1 TB
Maximum file system size	1 TB	1 TB
Number of Inodes	Dynamic, limited by disk space	Fixed, set at file system creation
Directory organization	B-tree	Linear

13.1.8 Reliability, Availability, and Serviceability (RAS)

This section discusses the RAS features of AIX 5L.

Error log scalability enhancements

Enhancements are made to the error log to detect consecutive duplicate errors and help prevent the error log from being overloaded. These enhancements provide a mechanism to keep track of the number of times an exactly identical error at the bit level occurs within a limited time period. An entry will be added to the error log along with the number of occurrences for this error. Examples of this type of error include:

- ▶ Floppy drive not ready
- ▶ External drive off line
- ▶ Ethernet card unplugged
- ▶ Defective line

This function is automatically enabled.

System hang detection

The `shdaemon` command function aids system administrators by providing a SMIT-configurable mechanism to detect system hangs and initiate the configured action. One of the actions in Table 13-7 on page 540 can take place when a system hang occurs.

Table 13-7 Actions at system hang condition

Actions	Default value
Error messages are displayed to notify system administrators.	off
Error messages are logged in the error report.	off
A high priority login is created so that system administrator can log in and perform problem determination.	on
Specific commands are executed according to the configuration setup.	off
Systems are rebooted based on the system configuration setup.	off

System administrators can set up these actions with the appropriate priority and turn them on or off as needed. An additional benefit of the tool is that it can be set with a threshold low enough that when that level occurs, a log entry is generated. The frequency of log entries can provide a measure of saturation and resource utilization.

Error log retrieval API for diagnostics

As part of the error log analysis performed by diagnostics, the `errpt -g` command is invoked to get raw data for each error log entry. This data is parsed and then put into a data structure for use by the calling application.

AIX 5L Version 5.1 provides the calling application with an application programming interface (API) to get data directly from the error log. This new method for obtaining data from applications is much faster when compared to the `errpt raw` mode. With this enhancement, error reports will contain both the error report and the diagnostics analysis. Previously, the error report only showed the error. A separate command was required to see the diagnostic analysis.

With this change, the error report will include diagnostic analysis for errors that have been processed. This may include such errors as:

- ▶ Disk errors
- ▶ Network controller errors
- ▶ Other hardware errors

Diagnostic analysis will give an indication of whether the hardware needs replacing or if the error may be safely ignored.

Automatic dump analysis

This tool enables accelerated customer support with less customer interaction time. The automatic dump analysis tool is capable of automatically examining a dump and pulling out the relevant information from a text for forwarding to support entities as an e-mail attachment. This will allow early diagnosis of dumps without having to send the entire dump file.

Dump reliability improvements

System dump capability has been improved to increase reliability. These improvements will help ensure that system dumps are consistent and reliable for easy problem determination. As dumps approach 3 GB in an uncompressed state, situations arise where dumps are incomplete or unavailable to technical support resources. This improvement adds compression as soon as possible, adds a cron job to the estimated dump size, and sends an error message recommending a larger dump device when it is appropriate. For systems that have paging space as the default dump device, this enhancement helps avoid the problem of insufficient paging space for dumps. The **snap** command is modified to eliminate an additional copy of the dump when it gathers files for its image. This reduces the space required to take a snap image.

Dedicated dump device

In AIX Version 4.3.3 and earlier, the paging space is used as the default dump device created at installation time. AIX 5L Version 5.1 servers with a real memory size larger than 4 GB will have a dedicated dump device created at installation time. This dump device is automatically created and no user intervention is required. The default name of the dump device is `lg_dumplv`. This name and the size of the dump device can be changed by using the `bosinst.data` file on a diskette at boot time.

Generate core files

AIX 5L Version 5.1 provides a capability to generate a core file for an application without requiring application termination. This capability helps increase the application availability as well as serviceability.

Kernel Debugger (KDB) Enhancements

The KDB kernel debugger and **kdb** command are enhanced. For AIX 5L and subsequent releases, the KDB kernel debugger is the standard kernel debugger and is included in the `unix_up`, `unix_mp`, and `unix_64` kernels, which may be found in `/usr/lib/boot`. KDB improvements include the following options:

- ▶ Disable paging of long outputs
- ▶ Print threads and processes in a long format when the entire table is printed
- ▶ Use command aliases

- ▶ Have multiple commands on a single line
- ▶ Use a command to display linked lists

13.1.9 Development and performance tools

The following sections provide a discussion of a select number of development and performance tools.

AIX Developer Kit JAVA2 Technology Edition Version 1.3.0

IBM AIX Developer Kit JAVA 2 Standard Edition Version 1.3.0 comes with a host of enhancements to Java classes and APIs, including the user interface, graphics, sound, networking, and math libraries. Of course, IBM's implementations, which are fully compliant with J2SE 1.3, have these enhancements too. IBM AIX Developer Kit Java 2 Technology Edition Version 1.3.0 has been engineered with the following features to deliver high performance and scalability to the most demanding e-business applications:

- ▶ Fully compatible with the Sun Java 2 Version 1.3.0 language, enabling write once run anywhere
- ▶ The latest version of the optimizing IBM Just-In-Time (JIT) compiler
- ▶ Efficient exploitation of AIX native threads
- ▶ Handle-less object model
- ▶ Fast and lightweight monitors
- ▶ Efficient management of large Java heaps through optimized object allocation and efficient garbage collection
- ▶ Thread-local heap
- ▶ Robust network support for a large number of concurrent connections
- ▶ Better scaling support for large numbers of threads and large numbers of file handles
- ▶ Tuned class libraries, which enhance performance in important areas like character codepage conversions, which are heavily used in many common types of Java applications.

The Web site specifically for Java on AIX is:

<http://www.ibm.com/developerworks/java/jdk/aix/>

Java security enhancements

In AIX 5L Version 5.1, a Java security enhancement has been made, providing several new APIs. These APIs are used by the Tivoli Security Toolkit. The new

APIs allow you to develop more secure Java applications and are provided with the following new Java enhancements:

- ▶ **Certificate Management Protocol (CMP)**
CMP provides support to online interactions between Public Key Infrastructure (PKI) components. For a full description of CMP, refer to RFC2510 and 2511 for CRMF. These RFCs are available at:
<http://www.ietf.org/rfc.html>
- ▶ **Java Cryptography Extension (JCE)**
JCE provides a framework and implementations for encryption and key handling. For more information about JCE, visit:
<http://java.sun.com/products/jce>
- ▶ **Java Secure Sockets Extension (JSSE)**
Java Secure Sockets Extension (JSSE) enables secure Internet communications. It provides a Java version of Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. For more information about JSSE, visit:
<http://java.sun.com/products/jsse>
- ▶ **Public-Key Cryptography Standards (PKCS)**
IBM Public-Key Cryptography Standards (PKCS) implementation supports the following RSA standards: PKCS #1, #3, #5, #6, #7, #8, #9, #10, and #12. For more information about PKCS, go to:
<http://www.rsasecurity.com/rsalabs/pkcs/index.html>

Performance analysis tools enhancements

For more accurate system performance analysis with more performance information, performance analysis tools include the following new tools:

- ▶ **truss** allows the tracing of all system calls made and signals received by a command or an existing process.
- ▶ **alstat** is a new tool, which reports alignment exception statistics. This tool can be used to detect performance degradations caused by misaligned data or code.
- ▶ **emstat** is a new tool for reporting emulation exception statistics. It can be used to detect performance degradation caused by emulated code, for example, code compiled for older systems and running on the latest PowerPC processors.
- ▶ **locktrace** permits dynamic enablement of kernel lock activity tracing. If the **bosboot -L** command was previously executed and the machine rebooted,

activity tracing could be turned on at the class level to minimize the overhead of lock tracing.

- ▶ **wlmon** permits analysis of workload manager class activity by resource. It is a Java-based Graphical User Interface for generating a trend report on the last 24-hours of WLM activity.
- ▶ The **perfstat** API is a new set of APIs, which provides easy access to kernel performance metrics.
- ▶ A new set of APIs is available to provide access to performance monitor data on selected processor types, namely 604, 604e, POWER3, POWER3-II, RS64-II, RS64-III, RS64-IV, POWER4, and POWER4+. Other processors of the POWER platform not listed are not supported by this API.

In addition, the Performance analysis tools include the following enhancements:

- ▶ **filemon** has been enhanced to process offline trace files to allow reports from the busiest systems to be processed as well.
- ▶ **gennames** is a newly documented utility, which supports offline analysis of trace files by the **filemon**, **netpmon**, **pprof**, and **tprof** analysis tools.
- ▶ **iostat** now provides adapter and system level throughput statistics.
- ▶ **netpmon** now can process off-line trace files to improve usability and scalability on larger systems. This tool has also been enhanced to support all current adapters.
- ▶ **rmss** is updated to support larger memory systems and the 64-bit kernel.
- ▶ **svmon** has been enhanced to support Workload Management tier and superclass and subclass reports. Additionally, it has been updated to support the 64-bit kernel.
- ▶ **topas** has been updated with NFS and SMP statistics, support for Workload Manager classes, and new full screens for process and WLM views.
- ▶ **vmstat** has been enhanced with a new I/O view for displaying an alternative set of metrics focused on I/O activity.
- ▶ **tprof** now supports the profiling of Java applications through the Java Virtual Machine Performance Interface (JVMPi).

The tools have been withdrawn in AIX 5L Version 5.1 are **bf** (bigfoot), **bfrpt**, **stem**, **syscalls**, and **lockstat**.

Please refer to the manual pages for **svmon**, **truss**, and **locktrace** for similar functions supported by those tools.

Inventory Scout

Inventory Scout searches the user's system for microcode levels and compares those levels with an IBM database of latest levels. When used with Web-based Microcode Discovery Service, an HTML report is generated for the user and includes links to the latest levels of microcode for the system (except for SP systems).

Inventory Scout gathers Vital Product Information (VPD) from the user's machine and, when used with Web-based VPD Capture Service, uploads it to the IBM MRPD (Machine Reported Product Data) database. This VPD is useful in determining the correct components to ship when a Miscellaneous Equipment Specification (MES) upgrade is ordered.

Inventory Scout is covered more in detail later in Chapter 16, "Resource management" on page 685.

Additional administrative tools

Several new enhancements for administrators have been added:

- ▶ The `-d` option is an additional option of the `sar` command that provides useful statistics such as throughput and average queue depth. Many of these statistics were previously provided with the AIX `iostat` command. The new `-d` option is added to AIX 5L Version 5.1 for compatibility with other UNIX operating systems.
- ▶ The `restore` command is enhanced to maintain the non-sparseness of the database files, similar to the `tar` command. This command introduces a new flag, `-e`, to maintain the sparseness of database files after they are restored. The new command option successfully generates the file(s) as non-sparse on output if specified and sparse by default if not specified.
- ▶ The startup and shutdown utilities have been expanded to now log their activities in much more detail when logging is enabled. This allows easier service and maintenance.

13.1.10 Linux affinity

AIX 5L incorporates a strong Linux affinity through the AIX toolbox for Linux applications and the integration of the Linux development environment into AIX libraries. This makes it possible to compile and run Linux applications on AIX, providing the ideal background to support this fast growing and competitive market. Countless developers around the world are focused on developing applications for Linux systems, and now you can easily port these applications and run them directly on AIX, taking advantage of all the features and benefits this operating system offers. We will discuss Linux affinity in 13.4.3, "Linux affinity on AIX 5L" on page 602.

13.1.11 Network

This section will describe major enhancements in the networking area that AIX 5L Version 5.1 provides.

IP multi-routing and multiple gateway

Multi-path routing is added to provide users the ability to specify multiple routes to a destination. This allows system administrators to:

- ▶ Choose the option of either configuring multiple routes for load balancing or setting up alternative paths to direct network traffic when the best route can no longer perform its tasks
- ▶ Define multiple default gateways

Also, this increases network availability by:

- ▶ Detecting the failure of a next-hop gateway that is listed in its route cache
- ▶ Routing data through alternative gateways that are specified

Web serving enhancements

A set of APIs provided by Fast Response Cache Architecture (FRCA) enables e-business applications to cache data, such as Web content, in the Network Buffer Cache (NBC). Through these APIs, Web serving can be done in the kernel using FRCA and NBC, significantly reducing the path length and increasing the performance of e-business applications for net commerce. FRCA and NBC are enhanced from managing static data to managing dynamically generated data. FRCA will support the Hypertext Transport Protocol (HTTP) Version 1.1, a standard of IETF RFC, including the persistent connection aspect of that protocol.

Dynamic Feedback Protocol (DFP) support

The Dynamic Feedback Protocol provides a mechanism for reporting statistics to server load balancing (SLB) devices, for example, Cisco's Catalyst 4840G, Catalyst 6000, or LocalDirector, so that future connections can be handled by most available servers. DFP helps improve server response in a Multi-node Load Balanced (MNLB) environment.

Explicit Congestion Notification (ECN)

Enhancements were made to improve performance of TCP/IP over congested networks, especially for Web serving. ECN capability is only available on the TCP layer. For TCP, ECN has three new functions:

- ▶ Negotiation between the end points during connection set up to determine if they are both ECN-capable

- ▶ An ECN-Echo (ECE) flag in the TCP header, so that the data receiver can inform the data sender when a Congestion Experienced (CE) packet has been received
- ▶ A Congestion Window Reduced (CWR) flag in the TCP header, so that the data sender can inform the data receiver that the congestion window has been reduced

TCP splicing

TCP splicing helps push the data-relaying function of a proxy application (from server side socket to the client side socket or vice versa) into the kernel. Performance of proxy applications is improved by reducing the pathway length through which data must travel. Software products that are expected to take advantage of this function include Network Dispatcher and, potentially, the Web Traffic Express caching proxy server.

Virtual IP Address (VIPA)

In previous AIX releases, an application had to bind to a real network interface in order to get access to a network or network services. If the network became inaccessible or the network interface failed, the application's TCP/IP session was lost, and the application was no longer available.

To overcome application availability problems as described, AIX 5L offers support for virtual IP addresses (VIPA) for IPv4 and IPv6. With VIPA, the application is bound to a virtual IP address, not a real network interface that can fail. When a network or network interface failure is detected (using routing protocols or other schemes), a different network interface can be used by modifying the routing table. If the rerouting occurs fast enough, then TCP/IP sessions will not be lost.

Network interface backup mode

EtherChannel is a network aggregation technology and load balancing technology that allows you to produce a single large pipe by combining the bandwidth of multiple Ethernet adapters. In AIX 5L Version 5.1, the EtherChannel feature has been enhanced to support the detection of interface failures. This is called network interface backup.

In the network interface backup mode, the channel will only activate one adapter at a time. The intention is that the adapters are plugged into different Ethernet switches, each of which is capable of getting to any other machine on the subnet/network. When a problem is detected, either with the direct connection, or through inability to ping a machine, the channel will deactivate the current adapter and activate a backup adapter.

Note: The network interface backup feature is currently supported by 10/100 Ethernet FC 2968 (devices.pci.23100020.rte) and 4962 (devices.pci.1410FF01.rte) and gigabit Ethernet PCI card FC 2969 (devices.pci.14100401.rte). If you are using other devices, you may receive unexpected results.

Sendmail upgrade enhancements

AIX 5L Version 5.1 uses Sendmail Version 8.11.0. This version has several enhancements and changes:

- ▶ The sendmail files `sendmail.cf` and `aliases` have been moved to the `/etc/mail` directory. Links exist on the POWER platforms that are required for the migration to AIX 5L Version 5.1 from earlier releases of AIX.
- ▶ Sendmail supports the Berkeley DB 3.1.14 format to more efficiently store the `aliases.db` database file. Other databases used can store their data in the Berkeley database formats.
- ▶ Support for message submission agents.
- ▶ Multiple queues, memory-buffered pseudo files, and more control over resolver timeouts improve performance.
- ▶ The ability to connect to servers running on named sockets.
- ▶ Better LDAP integration and support for LDAP-based routing.
- ▶ Improved support for virtual hosting.
- ▶ Even better anti-spam control features.
- ▶ Several new map classes, which include `arith` and `macro`.

More information on Sendmail is available from the following Web site:

<http://www.sendmail.org>

Packet Capture Library

Previous AIX operating system releases and AIX 5L offer the Berkeley Packet Filter (BPF) as a packet capture system. In addition to that, AIX 5L introduces a Packet Capture Library (`libpcap.a`), which provides a high-level user interface to the BPF packet capture facility. The AIX 5L Packet Capture Library is implemented as part of the `libpcap` library, Version 0.4 from Lawrence Berkeley National Laboratory (LBNL).

The Packet Capture Library user-level subroutines interface with the existing BPF kernel extensions to allow users access for reading unprocessed network traffic. By using the new 24 subroutines of this library, users can write their own network-monitoring tools.

Quality of service (QoS)

A new method for regulating network traffic flows named Quality of Service (QoS) was introduced in AIX Version 4.3.3. The demand for QoS arises from applications such as digital media or real-time applications and the need to manage bandwidth resources for arbitrary administratively defined traffic classes.

AIX 5L further enhances the QoS implementation to support overlapping policies in the QoS manager. This is important when two or more overlapping policies are installed; policies can then be enforced in order of highest priority. To improve the manageability of a QoS configuration, AIX 5L also offers four new commands to add, delete, modify, and list QoS policies.

Token ring support for MPOA

AIX 5L Version 5.1 provides support for token ring for multi protocol over ATM (MPOA). Token Ring emulation support for MPOA 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.
- ▶ Data paths are reduced from many hops between routers to a single hop between end clients.

This support provides better interoperability between an ATM network utilizing AIX MPOA support and Token Ring networks.

IP fragmentation

The multi protocol over ATM (MPOA) implementation supports IPv4 without options. In AIX 5L Version 5.1, MPOA has been enhanced to support IP fragmentation.

MPOA IP packet fragmentation support provides for instances where there is a heterogeneous network environment with varying Multiple Transmit Units (MTUs). This capability extends to circumstances where there are network configurations and normal MTU path discovery is not available.

Virtual Local Area Network

Virtual Local Area Networks (VLANs) can be thought of as logical broadcast domains. A VLAN splits up groups of network users on a real physical network into segments of logical networks. This implementation supports the IEEE 802.1Q VLAN tagging standard, with the capability to support multiple VLAN IDs running on Ethernet adapters.

The IEEE 802.1Q VLAN support can be configured over any supported Ethernet adapters. If connecting to a switch, the switch must support IEEE 802.1Q VLAN. When configuring a VLAN network, ensure that all virtual adapters within the virtual network have the same VLAN ID. Each VLAN can have a different maximum transmission unit (MTU) value, even if sharing a single physical Ethernet adapter.

VLAN support is managed through SMIT. Type the `smit vlan` fast path from the command line and make your selection from the main VLAN menu. Online help is available. After you have configured a VLAN, configure the IP interface (for example, en1) for standard Ethernet or et1 for IEEE 802.3, using Web-based System Manager, SMIT, or the command line interface.

13.1.12 AIX system security

In this section, we will take a look at AIX system security features.

IBM SecureWay Directory Version 3.2

Version 3.2 of the IBM SecureWay Directory implements the Lightweight Directory Access Protocol (LDAP) Version 3.2 and is offered with the AIX operating system product at no additional charge. LDAP consists of two major functions: the client and the server.

The IBM SecureWay Directory Version 3.2 consists of the following components:

- ▶ `slapd`: The server executable
- ▶ Command line import/export utilities
- ▶ A server administration tool with a Web-browser based interface for configuration and administration of the directory
- ▶ A Java-based directory content management tool and online user guide
- ▶ Online administration help
- ▶ Online LDAP programming references (C, server plug-ins, and Java/JNDI)
- ▶ SecureWay Directory Client Software Development Kit (SDK) that includes C run-time libraries and Java classes

Some of the more significant enhancements and new functions and features include:

- ▶ Fine-grained Access Control.
 - In addition to the current support of grant or deny access to a specific directory object or an entire directory sub tree, enhancement includes:
 - Ability to set Access Control Information (ACI) for a specific attribute

- Ability to allow/deny access to which subjects (using entryOwner) can define the ACIs
 - Support IETF draft-ietf-ldapext-acl-model-04.txt for ACL model
- ▶ Performance Enhancement.
- Add a true backup/restore directory content, including schema definitions and server configuration, by using DB2 utilities to enhance the performance of backup/restore. Existing import/export data to/from LDIF format is still supported for heterogeneous directory interoperability.
 - A new connection model has been implemented, which enhances the scalability with a pool of threads serving a large number of connections (configurable). Performance is also enhanced by reducing connection time as a list of active threads are maintained to service subsequent connection requests.
 - Fast Server startup eliminates the first-time, slow-server startup (after it is configured) through the use of the 'First Touch' mechanism in such a way that the attribute table is not created until an entry is added that uses that attribute.
- ▶ Event Notification Support supports the application notification of directory events.
- It can be used either for client caching functions or for any other management applications that require a notification of directory event unfold.
 - This is done by client registration to the directory event options, such as 'Add', 'Delete', 'Modify', and 'ModifyRDN', or any other combinations. The Server will notify the client applications whenever an event of interest is at or below the applications DN.
- ▶ Security Audit Support
- Support of Directory Audit service improves the security of directory server.
 - System administrators can use this log facility to examine any suspicious pattern of activity in an attempt to detect a security violation. With time stamp and BindDN recorded in the audit log, the violator can be easily traced and detected.
 - Audit plug-in support allows any application to receive the audit data and filter to incorporate with other audit information. This feature allows the directory audit data to be a part of the centralized audit facility where the enterprise application is based.

- ▶ Transaction Support.
 - With LDAP extended support, a set of LDAP operations are performed as a unit of work for commit or rollback.
 - This unit of work is committed to the database only when the EndTransaction-Commit is returned or it is rolled to the prior operation state. This transaction support is limited to a single connection to a single LDAP server. A list of LDAP operations that forms a unit of work should be limited to the moderation (less than 100) number.
- ▶ Kerberos Version 5 Support.
 - Enhanced LDAP operation adds Kerberos Version 5 and additional authentication mechanisms along with existing Secure Socket Layer (SSL) Version 3-based authentication using x509v3 public key certificates.
 - The Kerberos authentication option is used for not only user authentication, but also used for the authentication option for Server replication.
 - Kerberos-based authentication enables LDAP applications and AIX 5L Version 5.1 users to participate in a Single Sign-On environment within the Kerberos realm. This enhancement will enforce network security by not transporting the password on the wire.
 - Kerberos authentication is used for authenticated referrals. This helps to secure/validate the LDAP referral operation by using Kerberos-based authentication.
 - By using Kerberos cross-realm authentication support, the LDAP application can establish a transient trust with already established Kerberos authentication.

AIX LDAP security audit plug-in

Since the default audit function provided by the IBM SecureWay Directory may not be suited for the needs of the AIX security information management, an LDAP security plug-in has been added to AIX 5L Version 5.1. The LDAP security audit plug-in provides auditing of the LDAP security information server under the framework of the AIX security audit subsystem. The audit data is written into the AIX audit trails. With the AIX auditing commands, you can:

- ▶ Turn on/off auditing of the server
- ▶ Select the LDAP operations for auditing
- ▶ Manipulate the audit trails with the AIX auditing commands

The AIX LDAP Security Audit Plug-in provides administrators, who are familiar with the AIX system audit, a powerful tool to audit the IBM SecureWay Directory

server. The server auditing trails can be easily combined with other AIX audit trails to help administrators to analyze many events.

Native Kerberos Version 5 KDC Server/Client Support

Network Authentication Service is the IBM implementation of the MIT Kerberos Version 5 Release 1.1 Network Authentication Service. Network Authentication Service negotiates authentication and optionally encrypted communication between two points on the Internet or between components on a system.

Services provided will include the Generic Security Service Application Programming Interfaces (GSS-API) Version 2 and the key distribution server, `kadm5`.

A system can be configured to use Network Authentication Service for user password authentication, which will potentially make AIX 5L Version 5.1 authentication interoperative with other systems that use the MIT Kerberos Version 5 based authentication.

Services are provided to allow Kerberos to be configured as the systems default authentication mechanism. When AIX 5L Version 5.1 is configured to use Kerberos as the default authentication mechanism, all services that use the authentic routine to gain system access will then go to the Kerberos server, key distribution center, to seek user authentication. The users successfully logged in will have a Ticket Granting Ticket (TGT). There will also be a service provided to migrate existing users to Kerberos authenticated users.

Pluggable Authentication Mechanism

Pluggable Authentication Mechanism (PAM) is a flexible mechanism for authenticating users. This feature is added to support the need for a native AIX implementation of the X Single Sign-On (XSSO)/Pluggable Authentication Mechanism (PAM). This stand-alone module library, `libpam.a`, is located in `/usr/lib` directory. The PAM specification defines an API to an XSSO Sign-on Service for use by sign-on applications. The XSSO Service API is independent of the specific authentication mechanisms used. There are two distinct aspects to a Single Sign-on service:

- ▶ The Primary Sign-on operation, in which a user signs onto the policy domain as a whole
- ▶ The Secondary Sign-on operations, in which a user signs onto a service within the domain.

The principal objective of a Single Sign-On service is that secondary sign-on operations can be transparent to the user.

13.1.13 Networking security

This section describes the networking security enhancement of AIX 5L.

PKCS support

AIX 5L Version 5.1 offers an implementation of the cryptographic API PKCS#11 Version 2.01. PKCS#11 is a industry standard for accessing cryptographic hardware devices. In addition, AIX 5L Version 5.1 offers support for IBM 4758 Model 2 cryptographic coprocessor under the operating system PKCS#11 shared object. The PKCS#11 implementation is enhanced to utilize future IBM cryptographic hardware devices through the same shared library.

Applications available to utilize PKCS#11 include the iPlanet server suite. For additional information on PKCS11, refer to the RSA Laboratories Web site at:

<http://www.rsasecurity.com/rsalabs/pkcs/pkcs-11/>

IP Key Encryption Security

The Internet Key Exchange protocol to provide Virtual Private Networking (VPN) support is enhanced to enable the use of Certificate Revocation Lists (CRL) when authenticating remote users or devices. This is an improvement to scalability of VPNs through the use of Digital Certificates for a large number of users. When CRLs are used, digital certificates provide credentials for authentication, and individual users may be revoked by specifying their certificate number to the CRL. This simplifies network management by allowing one policy to be defined at the server level, and verifies that the certificate is valid and not contained in the CRL. CRLs may be fetched through HTTP or LDAP using socks4 or socks5 protocol.

The Web-based System Manager user interface for setting up tunnels has been streamlined and simplified. A full-function wizard guides the user through initial IKE tunnel definition. Policy information has been reorganized to make IP Security tunnel configuration more intuitive and require fewer steps.

Other IKE enhancements include the use of the commit bit to synchronize the use of Security Associations, and the definition of default policies to simplify the configuration for networks using dynamic IP addresses or DHCP. System administrators can define a Virtual Private network by one policy and a list of group members. They can also define default policies to specify the security parameters that are to be used when the addresses are dynamically assigned.

IKE support has also been extended to include IP Version 6 protocols. The IP Security functions for AIX 5L Version 5.1 now include the definition of static filters for IP Versions 4 and 6, and manually and dynamically defined private tunnels using IP Security protocol over IP Version 4 and 6 networks. IKE enhancements

include VPN functionality enabling users to import IKE tunnel configurations between Linux and AIX platforms. The `ikeconvert` script will process a Linux configuration file into an XML format suitable for loading into AIX.

13.2 AIX 5L Version 5.2

AIX 5L moves the operating system into the next stage of IT infrastructure self-management with innovative scalability technology while continuing to offer application flexibility with Linux, tools that simplify systems management, leadership security mapping between heterogeneous platforms, and affinity with pSeries focus market segments. The addition of dynamic logical partitioning and keyed Capacity Upgrade on Demand significantly improves flexibility, workload management, and system resource use in the datacenter.

AIX 5L Version 5.2 reliability and scalability, combined with application binary compatibility across all AIX Version 5 releases and concurrent 32/64-bit functionality, make it the best choice for customers who:

- ▶ Require a robust platform for business-critical applications
- ▶ Want to leverage their IT investments in technology and skills
- ▶ Have network interoperability requirements with heterogeneous systems
- ▶ Need components and tools to build tailored solutions
- ▶ Want to reduce the cost of computing through improved systems and network management
- ▶ Need security at all levels of their operating, application, and network environments
- ▶ Deploy applications worldwide requiring multilingual support

In this section, we discuss the enhancements introduced in AIX 5L Version 5.2.

13.2.1 Hardware requirements

AIX 5L Version 5.2 exclusively supports PCI architecture machines. Support for Microchannel Bus Architecture (MCA), Personal Computer Memory Card International Association (PCMCIA), and Instrumentation Systems and Automation Society (ISA) devices has been withdrawn.

- ▶ Version 5.2 withdraws support for the following architectures:
 - MCA (built-in and plug-in)
 - PCMCIA (built-in and plug-in)
 - ISA (PReP built-in and plug-in, although CHRP built-in support remains)

- ISA (CHRP plug-in)
- ▶ Version 5.2 withdraws support for the following processors:
 - Power 1
 - Power 2
 - Power Single Chip (RSC)
 - Power 2 Single Chip (P2SC)
 - 601
 - 603
- ▶ Version 5.2 withdraws support for PReP-specific functions for the following packages:
 - PReP PAL®
 - PReP desktop power management (hibernate)
 - All IDE support
 - All plug-in ISA adapter support
 - All PReP built-in ISA adapter support (although support for CHRP built-in ISA support remains)
 - PReP boot image from AIX Install CD-ROM
 - PReP boot image from AIX Standalone Diagnostics CD-ROM

Version 5.2 withdraws support for selected PCI adapters that are only supported on PReP platforms, as provided in Table 13-8.

Table 13-8 Version 5.2 withdrawn PCI adapter support

Feature	Description
2408	10-95 F/W SCSI SE, PCI/SHORT/32BIT/5V
2409	10-95 F/W SCSI DIFF, EXT ONLY, PCI/SHORT/32BIT/5V
2638	04-97 VIDEO CAPTURE(NTSC/PAL/SECAM), PCI/LONG/32BIT/5V
2648	06-95 (GXT150P) PCI/SHORT/32BIT/5V, GRAPHICS ADAPTER
2657	10-95 S15 GRAPHICS ADPTR, PCI/SHORT/32BIT/5V, WEITEK P9100
2837	04-97 MVP MULTI-MONITOR ADPTR, PCI/LONG/32BIT/3.3 OR 5V
2839	GXT100P Graphics Adapter
2854	10-96 (GXT500P), PCI/LONG/32BIT/3.3 OR 5V, GRAPHICS ADAPTER

Feature	Description
2855	10-96 (GXT550P), PCI/LONG/32BIT/3.3 OR 5V, GRAPHICS ADAPTER
2856	06-95 PCI/SHORT/32BIT/3.3 OR 5V, 7250 ATTACH ADAPTER
7252	GXT1000, 7250-002 Internal Graphics Accelerator
7253	GXT1000, 7250-002 with graphics feature
7254	Video Output Option
8242	06-95 10/100BASET ETHERNET PCI/SHORT/32BIT/5V, (3COM)

Version 5.2 withdraws support for PReP-specific ISA adapters (plug-ins), as provided in Table 13-9.

Table 13-9 Version 5.2 withdrawn PReP-specific ISA adapter support

Feature	Description
2647	06-95 VIDEO CAPTURE ENHANCEMENT, ISA/SHORT
2701	10-95 4 PORT SDLC, ISA/LONG, EIA 232/V.35/X.21, (GALE)
2931	10-95 8-PORT, ISA/LONG, EIA232 ADPTR/FAN-OUT BOX
2932	04-96 8-PORT, ISA/LONG, EIA232/422 ADAPTER
2933	10-95 128-PORT, ISA/LONG, EIA232 ASYNCH CONTROLLER
2961	10-95 1 PORT X.25, SDLC, PPP, ISA/LONG, ADAPTER (C1X)
2971	06-95 TOKEN RING ADAPTER, ISA
2981	06-95 ETHERNET ADAPTER, ISA, RJ45/BNC
8240	06-95 A/M 3COM ETHERNET ISA/SHORT TP ONLY
8241	06-95 A/M 3COM ETHERNET ISA/SHORT BNC/AUI

Note: Often a CHRP package would pre-req or co-req a PreP package to pull in required files to in order for the package to work. These selected files have now been moved to the CHRP packages and so no longer have a dependency on the PReP package, which has been removed.

Version 5.2 withdraws support for the following ISA adapters (plug-ins), even though they may run on a pSeries machine that is supported by Version 5.2. These include, but are not limited to, the adapters identified in Table 13-10 on page 558.

Table 13-10 Version 5.2 withdrawn ISA adapter support

Feature	Description
2931	8-PORT, ISA/LONG, EIA232 ADPTR/FAN-OUT BOX
2933	128-PORT, ISA/LONG, EIA232 ASYNCHRONOUS CONTROLLER
2932	8-PORT, ISA/LONG, EIA232/422 ADPTR/FAN-OUT BOX
2961	1 PORT X.25, SDLC, PPP, ISA/LONG, ADAPTER
2701	4 PORT SDLC, ISA/LONG, EIA 232/V.35/X.21

AIX Version 4.3 removed support for all AIX notebooks. All remaining PReP notebook support has been withdrawn from Version 5.2.

CHRP power management support is withdrawn in Version 5.2.

All MCA support is withdrawn in Version 5.2. The primary packages and support include:

- ▶ MCA PAL
- ▶ All plug-in and built-in MCA support
- ▶ MCA boot image from AIX Install CD-ROM
- ▶ MCA boot image from AIX Standalone Diagnostics CD-ROM
- ▶ Pegasus and other MCA-specific commands

In some cases, a CHRP plug-in and built-in I/O package will prerequisite or corequisite an MCA package to pull in the required files. In all cases, the CHRP packages have been rebuilt to include the files that they require, thus removing any dependency on the MCA package. The MCA package has also been removed.

Version 5.2 withdraws support for all PCI RS/6000 systems based on the PReP architecture and corresponding features including, but not limited to, the following, noting that all notebook support was withdrawn with Version 4.3, as provided in Table 13-11.

Table 13-11 Version 5.2 PCI RS/6000 withdrawn support listing

Systems	Family	Systems	Family
7020-0U0	40P	6015-066	
7020-SPE	40P	7248-100	43P
7020-B1B	40P	7248-120	43P

Systems	Family	Systems	Family
7020-B1C	40P	7248-132	43P
7020-D1D	40P	7043-140	
7020-D2D	40P	7043-240	
7020-D4E	40P	7024-E20	
6042-850	Notebook	7024-E30	
7247-821	Notebook	7025-F30	
7247-822	Notebook	7025-F40	
7247-823	Notebook	7317-F3L	
7247-860	Notebook	7026-H10	
6050	All models	6070	All models

Version 5.2 withdraws support for all MCA RS/6000 models and corresponding features including, but not limited to, the machines listed in Table 13-12.

Table 13-12 Version 5.2 MCA RS/6000 withdrawn support listing

Systems			
7006-41T	7006-41W	7006-42T	7006-42W
7007-N40	7008-M20	7008-M2A	7009-C10
7009-C20	7010-120	7010-130	7010-140
7010-150	7010-160	7011-220	7011-22G
7011-22S	7011-22W	7011-230	7011-23E 230E
7011-23S	7011-23T	7011-23W	7011-250
7011-25E 250E	7011-25F 250FTURBO	7011-25S	7011-25T
7011-25W	7012-320	7012-32E 320E	7012-32H
7012-340	7012-34H	7012-350	7012-355
7012-360	7012-365	7012-36T 36T	7012-370
7012-375	7012-37T 37T	7012-380	7012-390
7012-397	7012-39H	7012-G02	7012-G30

Systems			
7012-G40	7013-520	7013-52H	7013-530
7013-53E 530E	7013-53H	7013-540	7013-550
7013-55E 550E	7013-55L	7013-55S 550S	7013-560 560
7013-56F 560F	7013-570	7013-57F 570F	7013-580
7013-58F 580F	7013-58H	7013-590	7013-591
7013-595	7013-59H	7013-J01	7013-J30
7013-J40	7013-J50	7015-930	7015-950
7015-95E 950E	7015-970	7015-97B	7015-97E 970E
7015-97F 970F	7015-980	7015-98B	7015-98E 980E
7015-98F 980F	7015-990	7015-99E	990E 7015-99F 990F
7015-99J 990J	7015-99K 990K	7015-R10	7015-R20
7015-R21	7015-R24	7015-R30	7015-R3U R30U
7015-R40	7015-R4U R40U	7015-R50	7015-R5U R50U
7030-397	7030UPGRD	7030-3AT	7030-3BT
7030-3CT	7202-900		

Version 5.2 withdraws support for MCA-based SP nodes to the machines listed in Table 13-13.

Table 13-13 Version 5.2 MCA-based SP nodes withdrawn support

Feature	Description
2001	62 MHz Thin Nodes
2002	66 MHz Thin Nodes
2003	66 MHz Wide Node
2004	66 MHz Thin Nodes
RPQ	66 MHz Wide (59H)
2005	77 MHz Wide Node
2006	112 MHz High Node

Feature	Description
2007	135 MHz Wide Node
2008	120 MHz Thin Nodes
2009	200 MHz High Node
2022	160 MHz Thin Nodes

Version 5.2 withdraws support for the devices listed in Table 13-14.

Table 13-14 Version 5.2 device support withdrawn

Feature		
7027-HSC	PDOG	SE
7027-HSD	PDOG	DIFF
7236-001	ADEC	DRWR
7317-D10	DSK	DRWR
7318-P10		
7318-S20		
7319-100		
7319-110		

13.2.2 Software requirement

Systems operating on AIX 5L for POWER Version 5.2 are supported only when used within the system operating environments described in the appropriate hardware announcements and when used within the specified operating environment. When systems operating on AIX 5L for POWER Version 5.2 are used with other software or software in later announcements, other limitations may be included.

Table 13-15 Software requirement

Software	Requirements
AIX 5L Version 5.2	<ul style="list-style-type: none"> ▶ At least 128 MB of physical memory ▶ At least 2.2 GB of disk size
CSM	APAR IY34493 and IY34724

HMC recovery software Version 1.3 for pSeries

New systems include preloaded HMC software and firmware levels to support dynamic LPAR and Capacity Upgrade on Demand. The firmware release information is available from:

<http://techsupport.services.ibm.com/server/mdownload>

The HMC Recovery Software for pSeries is available to order from the Software Delivery and Fulfillment.

IBM Directory Version 4.1

To install the IBM Directory, your computer must meet the following minimum system requirements:

▶ IBM Directory client:

For the latest information on supported versions of AIX, go to `/usr/ldap/readme/lang/readme/client.txt`, or, if you are using a Web browser, go to `/usr/ldap/web/lang/readme/client.htm`.

A minimum of 128 MB RAM is required, though 256 MB is strongly recommended.

▶ IBM Directory server (including the client):

For the latest information on supported versions of AIX, refer to the server README file in `/usr/ldap/readme/lang/readme/server.txt`, or, if you are using a Web browser, go to `/usr/ldap/web/lang/readme/server.htm`.

In addition to the client requirements, the server requires one of the following Web servers, or a later version, installed and configured:

- IBM HTTP Server 1.3.19
- Lotus Domino Enterprise 5.0.2b Web server
- Apache Server 1.3.12
- Netscape FastTrack Server 3.01
- Netscape Enterprise Server 3.6.3, 4.0

Note: You must have a secure Web server to ensure that the transmission of your administration data is secure. For instructions, refer to the online documentation.

DB2 Version 7.2 is included with the IBM Directory. Higher versions of DB2 might also be supported. If you already have DB2 installed, you need approximately 45 MB of disk space. You need approximately 135 MB of disk space for both the IBM Directory and DB2. Disk space required for data storage depends on the number and size of database entries.

13.2.3 AIX 5L Version 5.2 release and support date

Table 13-16 provides information of the release and withdrawal date of AIX 5L Version 5.2.

Table 13-16 Release and withdrawal date

Program Number	AIX Release level	Available date	Customer service withdrawal date
5765-E62	AIX 5L Version 5.2	Oct. 18, 2002	Sept. 30, 2006

13.2.4 AIX 5L Version 5.2 highlights

AIX 5L Version 5.2 offers:

- ▶ Dynamic logical partitioning for processors, memory, and I/O
- ▶ Dynamic Capacity Upgrade on Demand
- ▶ Enhancements to Scalability and Workload Manager
- ▶ Enhancements to Enterprise Storage Management
- ▶ Cluster Systems Management for monitoring and administering multiple machines (both AIX and Linux) from a single point of control
- ▶ Advanced RAS features
- ▶ Additional security features and enhancements
- ▶ Linux affinity
- ▶ Network enhancements including Mobile IPv6, SNMPv3, and upgrade to BINDv9
- ▶ APIs from the latest C language and single UNIX specification standards

13.2.5 System management

AIX 5L Version 5.2 provides many enhancements in the area of systems management. This section discusses these enhancements.

Cluster Systems Management (CSM)

CSM for AIX can manage multiple machines from one single point of control. This solution for distributed systems management allows a system administrator to install and set up a cluster that can include pSeries servers, xSeries servers, or both. The system administrator can control the cluster by using remote parallel command execution, file synchronization, hardware control, and distributed monitoring with automated responses.

CSM for AIX enables xSeries servers running CSM Version 1.3 for Linux to join the cluster and be managed along with pSeries servers from the same single point of control.

Use of CSM is authorized only under the terms of either 60-day Try and Buy or a full production license (5765-F67) and LUM key.

Alternate disk install enhancement

Alternate disk install migration for network installation management (NIM) is now configurable through both the command line and SMIT. It is also possible to install the software from the BOS installation menus at system install time.

Alternate disk install at BOS installation time

There are two ways to install the software. It is now possible to install alternate disk installation at BOS install time, and the usual way with the **installp** command. When installing a new system from the AIX CDs, it is possible to install the software necessary to use alternate disk installation once the system is fully operational.

The following filesets are required to install the software necessary to enable alternate disk install:

- ▶ bos.alt_disk_install.rte
- ▶ bos.alt_disk_install.boot_images

These filesets can be installed during the BOS install (by selecting from the menus), or later.

Enabling NIM alternate disk migration

With AIX 5L Version 5.2 is a NIM alternate disk migration option available through the **nimadm** command and a SMIT fast path (**smitty nimadm**). The **nimadm** command (network install manager alternate disk migration) is a utility that allows the system administrator to create a copy of rootvg to a free disk (or disks) and simultaneously migrate it to a new version or release level of AIX. **nimadm** uses NIM resources to perform this function.

There are several advantages to using **nimadm** over a conventional migration:

- ▶ Reduced downtime.
- ▶ **nimadm** facilitates quick recovery in the event of migration failure.
- ▶ **nimadm** allows a high degree of flexibility and customization in the migration process.

Comparison reports for LPPs

System administrators can maintain a proactive maintenance strategy using the comparison report feature. Users can compare the filesets installed on their systems with the contents of an image repository or a list of available updates that can be downloaded from the pSeries Support Web site. This function can be executed from the command line (**compare_report**) or SMIT (**smit ty compare_report**). Selective reports can be used as input to request updates from the Web site.

13.2.6 Resource management

We will discuss the resource management features of AIX 5L Version 5.2.

Dynamic LPAR

With the availability of the IBM @server pSeries 690 server in December 2001, static logical partitioning (LPAR) was introduced to the pSeries platform. While LPAR provides a solution to logically remove and assign resources from one partition to another, the operating system in all affected partitions has to be rebooted, and the partitions have to be reset.

Dynamic LPAR (DLPAR) on IBM @server pSeries servers enables the movement of hardware resources (such as processors, memory, and I/O slots) from one logical partition running an operating system instance to another partition without requiring reboots and resets. With DLPAR technology, the following features are enabled: Dynamic reconfiguration, Dynamic Capacity Upgrade on Demand (DCUoD), and CPU sparing.

DLPAR system is made up of several components. To provide the foundation for DLPAR, the following components were made DLPAR aware:

- ▶ HMC
- ▶ Hypervisor
- ▶ Global-Firmware
- ▶ Local-Firmware
- ▶ AIX

Dynamic Capacity Upgrade on Demand

Dynamic Capacity Upgrade on Demand enables you to order and install POWER4 systems with additional hardware resources (processors) than what is currently required and keep these resources in reserve until they are required as future business application workloads dictate. To enable the additional resources, the system administrator can dynamically turn on the resources using DLPAR services without having to bring down the system. Additional CPUs can be

enabled by invoking the **chcod** CUoD command. This command can only be run by the super user or a user with system group membership. This function is well suited for those who have unpredictable application workload requirements.

Dynamic CPU guard

In AIX 5L Version 5.2, the CPU Guard implementation has been changed and enhanced to work in the new DLPAR Framework. The actual deallocation of the CPU resource is performed in the DLPAR Framework by the dynamic CPU removal procedure.

The DLPAR mechanism allowing the dynamic processor removal is based on leaving holes in the logical CPU ID's sequence, unlike the former CPU Guard implementation, where holes in logical CPU IDs are not tolerated for compatibility reasons. The Dynamic Reconfiguration (DR) strategy is to abstract the status of the CPUs by having CPU bind IDs, which are a sequence of IDs 0 through N-1 representing only the on-line CPUs. This strategy provides better MCM-level affinity, thus breaking the assumption of uniform memory access from all CPUs by Runtime Predictive Deconfiguration of Processors (RPDP). With the DR approach, the load from the failing CPU is moved to a CPU that corresponds to the last CPU bind ID. Thus, the failing CPU bind ID and the last CPU bind ID are swapped, leaving a hole in the logical CPU ID sequence and making the last online CPU the failing processor. Therefore, the `bindprocessor` system call interface, the **bindprocessor** command, the **bindintcpu** command, and the `switch_cpu` kernel service have been changed to work with the CPU bind ID model instead of the logical CPU ID model.

CPU Guard dynamically removes a failing CPU, whereas CPU sparing replaces a CPU with a spare one under the cover. During the reconfiguration, no notifications of any kind are sent to the user, kernel extensions, or to user-mode applications that are CPU Guard or DR-aware.

Workload Manager (WLM) enhancements

WLM is designed to give the system administrator greater control over how the scheduler and Virtual Memory Manager (VMM) allocate CPU, physical memory, and I/O resources to processes. It can be used to prevent different jobs from interfering with each other and to allocate resources based on the requirements of different groups of users.

The major use of WLM is for large SMP systems, and it is typically used for server consolidation, where workloads from many different server systems are combined, such as print, database, general user, transaction processing systems, and so on. These workloads often compete for resources and have differing goals and service level agreements. At the same time, WLM can be used in uniprocessor workstations to improve responsiveness of interactive work

by reserving physical memory. WLM can also be used to manage individual SP nodes.

WLM offers the system administrator the ability to create different classes of service and specify attributes for those classes. The system administrator has the ability to classify jobs automatically into classes, based upon the user, group, or path name of the application.

AIX Version 5.2 introduces new features to Workload Manager that improve its ease of use and provide more control over resource usage. There are five new enhancements in Version 5.2 for Workload Manager (WLM). They include:

- ▶ Attribute value grouping
- ▶ Event notification
- ▶ Time-based configurations
- ▶ Limits on total resources in a class
- ▶ Increase in the limit to the number of user-defined Superclasses and Subclasses.

13.2.7 Storage management

Here we discuss the storage management features of AIX 5L Version 5.2.

Increased file descriptor limit

AIX 5L Version 5.2 increased the maximum number of open file descriptors per process from 32767 to 65534. This limit is defined as `OPEN_MAX` in the include file `/usr/include/sys/limits.h`, as shown below:

```
#define OPEN_MAX 65534 /*max num of files per process */
```

File size enhancement

The Enhanced Journaled File System, JFS2, increases file and file system sizes to 16 TB with the 64-bit kernel. The 32-bit kernel continues to support 1 TB. Larger file systems created under the 64-bit kernel can be mounted only with the 64-bit kernel.

Universal Disk Format (UDF)

AIX 5L Version 5.2 includes support for the Universal Disk Format (UDF) Version 1.50 file system for the pSeries and RS/6000-supported DVD-RAM and DVD-ROM optical drives.

With the UDF file system, you can back up and restore AIX files and backup images. You can also start the system and install from disks created using the

mkcd command with DVD-RAM media. Software vendors and other users can create DVD distribution media for programs and data. AIX creates and supports file names only in 8-bit OSTA compressed UNICODE format.

Common HBA API support

Upper-level software applications that operate or use Fibre Channel (FC) Host Bus Adapters (HBAs) require FC information (for example, WWN or attached LUNs) for Storage Area Network (SAN) management or other reasons. The FC information is not available from HBAs in a consistent manner across operating systems, vendors, platforms, and in some cases not at all. Implementations to obtain such information are HBA vendor specific, for example, specific drivers or OS-specific calls have to be utilized to get to this information. This results in long qualification times, difficult integration across platforms, and inconsistency between HBA vendors, making implementation of SAN applications tedious to develop for upper-level software applications.

The Common HBA API, which is an industry standard programming interface for accessing management information in FC HBAs, provides a consistent low-level standard interface that can be implemented across vendors. Developed through the Storage Networking Industry Association (SNIA), the HBA API has been overwhelmingly adopted by SAN vendors to help manage, monitor, and deploy storage area networks in an interoperable way. With AIX 5L Version 5.1 ML 5100-03 and AIX 5L Version 5.2, support for the Common HBA API Version 1.92 has been added, with the exception of the `HBA_GetEventBuffer()`.

The Common HBA API is implemented as a set of C programming language library functions, which allow access to low level, FC HBA information, and the OS mappings.

Unaligned I/O support in LVM

In AIX 5L Version 5.2, file systems and kernel extensions have no LVM restrictions to contend with for size and alignment of I/O requests from the LVM strategy routine. A file system or kernel extension can now issue a single large I/O to the LVM strategy layer instead of breaking this I/O up into many individual smaller I/Os. This now allows LVM to issue a single iodone to the layer above LVM when the I/O is complete. The enhanced journal file system (JFS2) and AIO I/O requests currently take advantage of this feature.

Advanced RAID support

LVM enables a RAID to increase the size of a logical unit number (LUN) up to 1 TB. After a `varyoffvg` and `varyonvg` cycle completes, the new disk size becomes effective and LVM can use the additional disk space.

JFS2 snapshot image

AIX 5L Version 5.2 introduces the JFS2 snapshot image. The JFS2 snapshot image gives a consistent block level image of a file system at a given point in time. The snapshot will stay stable even if the file system that the snapshot was taken from (referred to hereafter as the snappedFS) continues to change.

The snapshot can then be used to create a backup of the file system at the given point in time that the snapshot was taken. The snapshot also provides the capability to access files or directories as they were at the time of the snapshot.

AIX 5L Version 5.2 provides the following functionality for a snapshot image:

- ▶ Snapshot creation on a separate logical volume from the snappedFS.
- ▶ Read-only access to a snapshot through a mounted file system.
- ▶ Read-only access to a snappedFS while the snapshot is created.
- ▶ Snapshot information listing.
- ▶ Snapshot removal.
- ▶ Capability of multiple snapshots for a file system.
- ▶ Snapshots are persistent when snappedFS is mounted or unmounted, and not persistent if a system crash occurs.
- ▶ Backup support for backbyname and backbynode

13.2.8 Reliability, Availability, and Serviceability (RAS)

This section introduces the RAS enhancement of AIX 5L Version 5.2.

System dump facility enhancements

The system dump facility has been enhanced to allow greater functionality in component dump routines. There is also support for unlimited dump size to allow a dump routine to return unknown amounts of dump data. Prior to Version 5.2, individual components would use the `dmp_add` and `dmp_del` services to register and unregister data areas to be included in the system dump. The components were each required to allocate and pin their own buffer space during initialization. The master dump table only has a pointer to the component's dump routine and has no visibility to the actual size of the component's dump data. This prevents the system from obtaining an accurate dump size estimate. When a system dump is started and the component's dump routine is called, the component is required to return all the dump data in one array. The maximum number of `cdt_entries` for the 64-bit dump is approximately 21840. This is problematic when the system has to dump data for 30000 processes.

System hang detection

Hang conditions on I/O devices can be detected and recovered without bringing systems to a slow halt. Based on a user-defined time-out threshold, lost I/O is detected and the system alerts the user. Users can also configure the system for an automatic reboot.

Boot LED displays

AIX 5L Version 5.2 provides enhanced support for the front panel display. The boot scripts now display additional information on the second line of the front display panel to give more information of specific LED values. During bootup, some of the LEDs can be displayed for an extended period of time. An example of this would be the 551 code, which is the **varyonvg rootvg** command. The second line for specific LEDs shows whether the phase is complete or if there is an error.

Multipath I/O

AIX 5L Version 5.2 provides a new feature called Multipath I/O (MPIO) that allows for a single device (disk, lun) to have multiple paths through different adapters. These paths must reside within a single machine or logical partition of a machine. Multiple machines connected to the same device are considered as clustering and not as MPIO. There are three reasons for MPIO:

- ▶ Performance improvement
- ▶ Improved reliability and availability
- ▶ Easier administration

13.2.9 Development and performance tools

We will discuss development and performance tools of AIX 5L Version 5.2.

AIX Developer Kit Java 2 Technology Edition Version 1.3.1

AIX Developer Kit Java 2 Technology Edition Version 1.3.1 32-bit Version for POWER is included with the AIX base operating system. The 64-bit version is available on the AIX Version 5.2 Bonus Pack.

Performance analysis tools

Performance analysis tools include support for thread (**curt**), lock (**sp1at**), procs (**truss** and **proc** commands), and JFS2 (**filemon** and **fileplace**).

Xprofiler analysis tool

The X Windows-based profiler (Xprofiler) is now included with the AIX 5L Version 5.2 operating system. Xprofiler is a tool that allows you to analyze your parallel

and serial applications. It uses procedure profiling information to construct a graphical display of the functions in your application. The graphical user interface (GUI) gives you a general overview of your application and allows you to focus on CPU-intensive sections of your application.

AIX tuning framework

Prior to AIX 5L Version 5.2, all the performance parameters that can be set by the **vm tune**, **sched tune**, **no**, or **nfso** commands were lost at the next system reboot. The syntax and the output of those commands were also completely different. With AIX 5L Version 5.2, system administrators can use systems management interfaces (Web-based System Manager and SMIT) to manage performance tuning parameters and preserve those settings whenever the system is rebooted. These interfaces support operating system, memory, network, and NFS tuning parameters.

Perl 5.8

The version of the Perl language installed with AIX is updated to 5.8.

13.2.10 Linux affinity

AIX 5L affinity with Linux helps enable faster and less costly deployment of multi-platform, integrated solutions across AIX and Linux platforms. Many applications developed on and for Linux will run on AIX 5L with a simple recompilation of the source code. IBM provides, at no cost, an AIX toolbox for Linux applications, which is a collection of Open Source and GNU software commonly found with Linux distributions. Because the applications run on AIX, businesses can combine the flexibility of Linux with the advanced features of AIX, including advanced workload management, sophisticated systems management tools, scalability, and security. We will discuss more about AIX affinity with Linux in 13.4.3, “Linux affinity on AIX 5L” on page 602.

13.2.11 Network

In this section, we will discuss enhancements in the networking area that AIX 5L Version 5.2 provides.

SNMPv3

AIX 5L Version 5.2 now supports Simple Network Management Protocol (SNMP) Version 3. Prior to Version 5.2, AIX only supported SNMP Version 1. From Version 5.2 on, the new SNMP agent is a SNMP Version 1/Version 2c/Version 3 compatible agent. SNMP provides a powerful and flexible framework for message level security and access control. SNMPv3 increases security

robustness through both authentication and encryption. Users can securely configure and monitor a system remotely.

Mobile IPv6

Mobile devices that use Mobile IPv6 enable the devices to keep the same global address on supporting networks. Applications maintain transport and higher-layer connections while changing locations. It allows mobility across homogenous and heterogeneous networks.

EtherChannel backup adapter

AIX 5L Version 5.2 introduces the concept of configuring a backup adapter to the EtherChannel. The backup adapter's purpose is to take over the IP and MAC address of the channel in the event of a complete channel failure, which is constituted by the failure of all adapters defined to the channel. It is only possible to have one backup adapter configured per EtherChannel.

All adapters that constitute the EtherChannel must be connected to the same switch. Version 5.2 can protect against a switch failure, as it provides the capability for the backup to be connected to a different switch to the EtherChannel. Therefore, to guard against switch failure and introduce further resilience, it is recommended that the backup adapter is connected by a separate Ethernet switch to the EtherChannel. Until takeover, the backup adapter is idle.

The process is as follows:

- ▶ If all but one of the primary adapters fail, then no action is taken, as the primary objective is to keep the EtherChannel open.
- ▶ If all primary adapters fail, the backup adapter is checked to see if it is functioning. If the backup adapter is down, the primary adapters stay as the active channel. This is because it is more likely that one of the EtherChannel adapters will come back up before the single backup adapter.
- ▶ If the backup adapter is up and all the primary adapters fail, then failover starts. All the adapters in the EtherChannel are disabled, and take on the MAC and IP address of the backup adapter. The backup adapter takes on the MAC and IP of the EtherChannel. All adapters are then re-enabled.
- ▶ Gratuitous ARPs are sent to ensure that the MAC associated with the EtherChannel port is now mapped to the backup adapter port.
- ▶ When at least one of the adapters in the EtherChannel becomes available, the MAC and IP are swapped back to the EtherChannel following the same process as before.

Quality of Service (QoS) enhancements

Prior to AIX Version 5.2, the policy agent managed all the policy management information and used a socket to communicate with the kernel. If the policy agent was stopped or ended abnormally, QoS would stop functioning. In AIX Version 5.2, the policy management information is still managed in the policy agent, but the policy agent publishes the policy management information into the QoS manager in the kernel. Because the QoS manager has a copy of the policy management information in pinned memory, QoS will still function if the policy agent is not running.

13.2.12 AIX system security

This section introduces AIX system security.

Public Key Infrastructure enhancements

AIX 5L Version 5.2 provides its own Certificate Authentication Service, with the ability to authenticate users using X.509 Public Key Infrastructure (PKI) certificates and to associate certificates with processes as proof of a user's identity. It provides this capability through the Loadable Authentication Module Framework (LAMF), the same extensible AIX mechanism used to provide DCE, Kerberos, and other authentication mechanisms.

PKI is a comprehensive system of policies, processes, and technologies working together to allow users and applications to exchange information securely and confidentially. PKI uses pairs of asymmetric keys, provided by a trusted third party known as a CA (Certificate Authority), to encrypt and decrypt information. These digital signatures provide the following security services:

- ▶ Entity authentication
- ▶ Data confidentiality
- ▶ Data integrity
- ▶ Non-repudiation
- ▶ Privilege management

PAM enhancements

AIX 5L Version 5.2 security services has been integrated with the Pluggable Authentication Modules (PAM) framework. The PAM framework allows administrators to incorporate multiple authentication mechanisms into an existing system through the use of pluggable modules. Applications written using the PAM framework do not need to be modified to support new authentication methods or modules.

In Version 5.1, the PAM libraries and include files were supplied but were not integrated into the AIX Security Services. In Version 5.2, applications that use the PAM framework could call AIX Security Services, and applications that use the AIX security libraries could now call PAM modules.

Cryptographically secure random number

A new version of the cryptographically secure random number generator supports cryptographic applications that require high-quality, secure random numbers. This version replaces the existing UNIX random number generator functions as part of the UNIX standard. AIX Internet Key Exchange (IKE) uses this random generator approach to increase its security.

IBM Directory Server Version 4.1

AIX 5L Version 5.2 now includes IBM Directory Server Version 4.1. IBM Directory Server Version 4.1 provides a powerful Lightweight Directory Access Protocol (LDAP) server that uses the IBM DB2 Universal Database™ Version 7.2 engine for reliability. The IBM Directory Server is an integral part of the new directory enablement features announced in AIX 5L Version 5.2. AIX supports a Certificate Authentication Service with Public Key Infrastructure (PKI), which stores PKI certificates in LDAP. The AIX System V printing subsystem is now directory enabled, allowing printer and print queue configuration to be stored in LDAP. AIX supports LDAP authentication and storage of user and group security attributes into LDAP. Network information services (NIS) maps can now be stored and accessed in LDAP.

The IBM Directory Server without SSL support is packaged on the AIX 5L Version 5.2 product media. The IBM Directory Server with secure socket layer (SSL) support is included on the expansion pack media.

13.2.13 Networking security

Here we take a look at network security features that AIX 5L Version 5.2 provides.

Generic data management tunnel support

The AIX IKE now supports the creation of a generic data management tunnel, also known as a phase 2 tunnel. This feature is used mainly when an IPSEC endpoint is using dynamic host configuration protocol (DHCP) to assign IP addresses. Normally, data management tunnels are identified by their IP address; with DHCP, the endpoint address is dynamic. The generic data management tunnel will be used if a request was authenticated by phase 1 and an IP address is not specifically configured in the database.

The generic data management tunnel is not a real tunnel, but a tunnel definition that is used when an incoming data management message does not match any defined data management tunnels. Defining a generic data management tunnel is optional, and there can only be one generic data management tunnel per key management tunnel definition. It can only be used in the case where the AIX system is the responder.

Diffie-Hellman group 5 supported

The AIX IKE has now been enhanced to support Diffie-Hellman (DH) group 5. Prior releases of the AIX only supported DH groups 1 and 2. Diffie-Hellman key exchange is a public key cryptosystem where public values are exchanged to arrive at a symmetric key among the end entities. Symmetric keys derived using Diffie-Hellman group 5 are more secure than the currently implemented Diffie-Hellman groups 1 and 2. With IP security static filter rules, users can enter a description of up to 80 characters.

13.3 AIX 5L Expansion, Bonus, and Web Download Pack

The AIX Expansion Pack is a collection of extra software that extends the base operating system capabilities. The AIX Bonus Pack is a collection of additional software. The AIX Web Download Pack is a collection of additional software available for download from the Web. It complements the AIX operating system with the benefit of additional packaged software at no additional charge.

Typical releases may include: database software, development tools, software supporting e-business, interoperability support, browsers, Java, and Internet application development tools, network management utilities, and country specific security encryption.

AIX 5L Expansion Pack, Bonus Pack, or Web Download Pack is intended for customers who:

- ▶ Want to exploit Internet technologies
- ▶ Want to take advantage of the latest in Java development tools
- ▶ Need security enhancements at all levels of their operating, application, and network support
- ▶ Need help in reducing the cost of computing through improved systems and network management
- ▶ Want to evaluate products that address PC interoperability requirements
- ▶ Need to view AIX documentation in HTML format

Note: This section of AIX 5L Expansion Pack, Bonus Pack, and Web Download Pack is based on information announced at the time of the writing of this redbook. As IBM continues to review and develop it, the contents of these packs can be changed or upgraded when this redbook is published.

13.3.1 System management

The following section describes the system management utilities included on the AIX 5L Expansion, Bonus Pack, and Web Download Pack.

IBM Network Authentication Service Version 1.3

IBM Network Authentication Service Version 1.3 is a network authentication protocol based on the IETF RFC 1510 standards protocol for the Kerberos Version 5 Network Authentication Service. It includes the Generic Security Service Application Program interfaces (GSS-API) and the key distribution center (KDC) server, which allows AIX middleware and external application writers to use authenticated message flow between their respective components. It provides enabled client libraries, headers, and a toolkit for building Kerberos based 64-bit applications. Triple DES encryption, LDAP directory support, KDC multi-thread, enhanced client support, and improved performance for concurrent login and Network Authentication Service server availability are also provided.

Additional product information can be found at:

<http://web.mit.edu/kerberos/www>

Certificate Authentication Service Version 5.2.0.10

Certificate Authentication Services Version 5.2.0.10 performs account authentication using Public Key Infrastructure (PKI) certificates.

Certificate Authentication Services Version 5.2.0.10 now offers:

- ▶ Support for IBM Grid Toolbox for AIX
- ▶ Installation and configuration usability enhancements
- ▶ Currency for X.509 standard
- ▶ Support for PKCS10

Cryptographic Library Version 5.2

Cryptographic Library Version 5.2 contains various new, up-to-date cryptographic functions, and includes an application programming interface.

Reliable Scalable Cluster Technology (RSCT) Version 2.3

Reliable Scalable Cluster Technology (RSCT) Version 2.3 adds data confidentiality to the RSCT subsystem. A range of cryptographic technologies are supported.

Encryption support for IBM Directory Version 4.1

Encryption Support For IBM Directory Version 4.1 includes:

- ▶ IBM Directory Server and Client Utilities for Maximum Encryption Version 4.1, which provides 128-bit and Triple DES encryption and includes a Java Naming and Directory Interface (JNDI) Client Software Developer Toolkit (SDK), which supports encryption up to 128-bit.

Additional product information can be found at:

<http://www.ibm.com/software/network/directory/server>

Encryption support for SecureWay Directory Version 3.2.2

Encryption Support For SecureWay Directory Version 3.2.2 includes:

- ▶ SecureWay Directory Server and Client Utilities for Maximum Encryption Version 3.2.2, which provides 128-bit and Triple DES encryption and includes a Java Naming and Directory Interface (JNDI) Client Software Developer Toolkit (SDK) which supports encryption up to 128-bit.

Additional product information can be found at:

<http://www.ibm.com/software/network/directory/>

IBM Global Security Toolkit Version 4.0

IBM Global Security Toolkit (Gskit) Version 4.0 provides certificate support for AIX 5L Internet Protocol.

Additional product information may be found at:

<http://www.ibm.com/software/network/directory/library/>

IBM Global Security Toolkit Version 5.0

IBM Global Security Toolkit (Gskit) Version 5.0 provides secure socket layer (SSL) support with Triple DES encryption for SecureWay Directory Version 3.2.2 and IBM Directory Server Version 4.1. It also provides certificate support for AIX 5L Internet Protocol and for SecureWay Directory Version 3.2.2 and IBM Directory Server Version 4.1.

Additional product information may be found at:

<http://www.ibm.com/software/network/directory/library/>

Data Encryption Standard (DES) library routines

Data Encryption Standard (DES) Library Routines for AIX Version 5.1 and Version 5.2 provide the capability for Electronic Code Book (ECB) and Cipher Block Chaining (CBC) encryption using 64-bit encryption keys.

IBM IP Security Version 5.1 and Version 5.2

IBM IP Security Version 5.1 and Version 5.2 provides encryption for 56-bit, and Triple DES support for the AIX 5L IP security enhancements.

Web-based System Manager Security Version 5.1 and Version 5.2

Web-based System Manager Security Version 5.1.0.50 and Version 5.2.0.10 provides 128-bit encryption, and it 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:

<http://www.ibm.com/servers/aix/wsm/>

Additional product information on AIX security can be found at:

<http://www.ibm.com/servers/aix/os/52desc.html>

Open Secure Shell (OpenSSH) Version 3.6

Open Secure Shell (OpenSSH) Version 3.6 is a set of client and server connectivity tools that encrypt, with Triple DES and Blowfish encryption, network traffic to help eliminate network level attacks. It provides NLS enablement and code, which contains support for user privilege separation.

Open Secure Shell Version 3.6.0.0 for AIX 5L Version 5.1 and Version 3.6.0.52 for AIX 5L Version 5.2 includes Kerberos Version 5 support compiled into the binaries. Kerberos is an optional authentication method for OpenSSH that is configured by the system administrator. If Kerberos is configured, the IBM Network Authentication Service Version 1.3 for AIX needs to be installed and configured. The AIX 5L Version 5.1 and Version 5.2 versions of OpenSSH have been updated using the openssh-3.6.1p1 level of source code from openssh.org. Additional product information can be found at:

<http://www.openssh.org>

SNMPv3 Encryption Version 5.2

SNMPv3 Encryption Version 5.2 provides CBC 56-bit DES encryption support.

13.3.2 e-business and e-commerce

We will discuss the e-business and e-commerce features included on the AIX 5L Expansion, Bonus Pack, and Web Download Pack.

WebSphere Application Server

WebSphere Application Server Version 5, the foundation of the WebSphere software platform, helps maximize your company's return on investment while leveraging existing software assets and making your administrative resources more productive. WebSphere Application Server Version 5 Trial Program demonstrates how WebSphere Application Server provides the core software to meet the changing demands of dynamic e-businesses by offering:

- ▶ Full Java 2 Enterprise Edition (J2EE) Version 1.3 compatibility
- ▶ Industry-leading support for key Web services standards
- ▶ Native, enterprise-ready Java Messaging Server (JMS)
- ▶ Integrated development tools
- ▶ Unparalleled connectivity
- ▶ Superior performance and scalability.

Additional product information can be found at:

<http://www.ibm.com/software/webservers/appserv>

IBM Directory Server Version 4.1

IBM Directory Server Version 4.1 provides a powerful Lightweight Directory Access Protocol (LDAP) identity infrastructure that is the foundation for deploying comprehensive identity management applications and advanced software architectures like Web services.

Additional product information can be found at:

<http://www.Ibm.com/software/network/directory/server>

IBM HTTP Server Version 1.3.19.4

IBM has enhanced the HTTP Server with performance and SSL for secure transactions. This product includes software developed by the Apache Group for use in the Apache HTTP server project. For more information on the project, refer to the Web site at:

<http://www.apache.org/>

Additional product information can be found at:

<http://www.ibm.com/software/webservers/httpservers/>

Wireless Application Protocol (WAP) Gateway Version 1.1

Wireless Application Protocol (WAP) Gateway Version 1.1 facilitates communications between the gateway and WAP-enabled devices using the WAP protocol.

Additional product information may be found at:

<http://www.kannel.3glab.org/>

Netscape Communicator 4.8

Netscape Communicator 4.8 provides 128-bit encryption with National Language Support. Bi-directional support for viewing Hebrew and Arabic HTML pages is also provided. Additional components to Netscape Navigator include Netscape Messenger (Netscape Mail), Collabra (Netscape discussion groups), and Composer (Netscape Web page publishing).

Additional product information can be found at:

<http://wp.netscape.com/browsers/4/index.html>
<http://www.ibm.com/servers/aix/browsers/index.html>

13.3.3 Development

The following sections discuss the development tools located on the AIX 5L Expansion, Bonus Pack, and Web Download Pack.

IBM AIX Developer Kit Java 2 Technology Version 1.3.1

IBM AIX Developer Kit Java 2 Technology Version 1.3.1 64-bit version for POWER and IBM AIX Developer Kit Java 2 Technology Edition Version 1.3.1 32-bit version for POWER are based on the Java 2 SDK 1.3 reference port and includes IBM functional and performance enhancements. The Developer Kit features IBM Just in Time Compiler (JIT) and Mixed Mode Interpreter. The 64-bit version allows customers to create and run 64-bit Java applications, improving overall application scalability. These programs also include tools to build secure java applications.

Additional product information can be found at:

<http://www.ibm.com/developerworks/java/jdk/aix/index.html>

Tools to build secure Java applications

Tools to build secure Java applications are provided for the IBM AIX Developer Kit Java 2 Technology Edition Version 1.3. These tools include:

- ▶ Java Cryptography Extension (JCE) Version 1.2.1
- ▶ Certification Management Protocol (CMP)

- ▶ Public Key Cryptography Standards (PKCS)
- ▶ Secure Multi-Purpose Internet Mail Extensions (S/MIME)
- ▶ Java Secure Sockets Extension Version 1.0
- ▶ Java Authentication and Authorization Service (JAAS) Version 1.0

Additional information may be found at:

<http://www.java.sun.com/products/jce>
<http://www.ietf.org/rfc/rfc2510.txt> (or 2511.txt or 2311.txt)
<http://www.rsasecurity.com/rsalabs/pkcs/>
<http://www.java.sun.com/products/jsse/>
<http://www.java.sun.com/products/jaas/>

Geodesic Systems Great Circle Version 6.0.1.5

Geodesic Systems Great Circle Version 6.0.1.5 (30-day evaluation software) is an advanced testing and diagnostic tool that allows software engineers to identify and resolve performance and reliability problems in C and C++ application software. It pinpoints elusive errors that can cause excessive application memory use and devastating crashes. Great Circle provides an intuitive and flexible Web browser-based user interface that supports remote monitoring and diagnosis of distributed applications, enabling software engineers to validate application performance in real-world environments.

AIX Toolbox for Linux application

This comes with a separate set of CDs that contain a collection of open source and GNU software built for AIX Version 4.3.3 and AIX 5L for IBM @server pSeries systems and IBM RS/6000. These tools provide the basis of the development environment of choice for many Linux application developers. All the tools are packaged using the easy to install RPM format. There is a strong affinity between Linux and AIX for applications. AIX has a long history of standards compliance and it is generally straight-forward to rebuild Linux applications for AIX. The AIX toolbox for Linux applications demonstrates the strong affinity between Linux and AIX. Additional information on tools for AIX application deployment may be found at:

<http://www.ibm.com/servers/aix/products/aixos/linux>

Modular I/O (MIO) Library Version 2.3

Modular I/O (MIO) Library Version 2.3 allows users to analyze the I/O of their application and then tune the I/O at the application level for a more optimal performance for the configuration of the current operating system. It includes a shared object for 64-bit executables.

Additional information may be found at:

http://www.research.ibm.com/actc/Opt_Lib/mio/mio_doc.htm

PartnerWorld for developers

PartnerWorld® for developers and tools for AIX application deployment Web sites provide tools and information to meet your requirements for developing applications and managing your production requirements. Solutions provided by IBM and IBM Business Partners are available to support the full spectrum of the life cycle of an application. All stages of application life cycle management are covered, providing an organized approach to assist you in locating the right tools for your application development and deployment needs.

Additional information on tools for AIX application deployment may be found at:

<http://www.developer.ibm.com/tech/>

13.3.4 Multimedia

As a multimedia utility, Adobe Acrobat Reader 5.06 allows you to view and print Portable Document Format (PDF) files. The Asian font packs for Acrobat Reader allow you to display PDF files that contain text in Chinese Simplified, Chinese Traditional, Japanese, and Korean.

Additional product information can be found at:

<http://www.adobe.com/products/acrobat/readermain.html>

Additional font pack information can be found at:

<http://www.adobe.com/products/acrobat/acrrasianfontpack.html>

13.3.5 Interoperability

AIX Fast Connect Version 3.1.2 (60-day evaluation software) is a full-function evaluation copy of AIX Fast Connect for AIX 5L provides file and print server capabilities to Windows 95, Windows 98, Windows NT, Windows 2000, Windows XP, and OS/2® servers and clients. This feature takes advantage of the scalability and performance capabilities of AIX. Support for Microsoft Distributed File System, authentication using LDAP and Level II Oplocks is provided. An unlimited number of clients are supported.

Additional product information can be found at:

http://www.ibm.com/servers/aix/products/ibmsw/system_man/fastconn.html

13.3.6 Storage management

In this section, we discuss the storage management solutions included on the AIX Expansion, Bonus Pack, and Web Download Pack.

IBM Tivoli Storage Manager Version 5.1.5

This 60 days evaluation software is a multi-function storage software product that addresses the challenges of complex storage management across distributed environments. It protects and manages a broad range of data and provides:

- ▶ High-speed automated data recovery
- ▶ Improved application availability during backups with online image backups, server-free backup and restore, and support for HACMP for AIX clients
- ▶ Improved throughput and reduced CPU utilization on AIX servers
- ▶ Compatibility with hundreds of storage devices, as well as LANs and WANs

Additional product information can be found at:

<http://www.tivoli.com/products/index/storage-mgr/>

IBM Tivoli Storage Resource Manager Version 1.2

Tivoli Storage Resource Manager Version 1.2, Evaluation Software is a comprehensive storage resource management (SRM) solution for storage resources across the enterprise. It enables administrators to identify, manage, control, and predict storage usage. Tivoli Storage Resource Manager provides a true consolidated view of storage, combining file system utilization with application utilization, across operating systems and storage devices.

Analyzing the storage usage of applications provides insight beyond traditional SRM tools. Designed for efficiency and ease-of-use, Tivoli Storage Resource Manager uses a single agent per server to provide detailed information without high consumption of network bandwidth or CPU cycles.

VERITAS Foundation Suite

Here we discuss the VERITAS Foundation Suite (60 day evaluation software) and VERITAS Foundation Suite High Availability (60 day evaluation software). VERITAS Foundation Suite is a leading on-line storage management software for enterprise computing environments. The Foundation suite combines file system and disk management technology to ensure easy management of all storage, optimum performance, and maximum availability of essential data. VERITAS Foundation Suite Evaluation Software is composed of:

- ▶ VERITAS File System Version 3.4.2.2 Evaluation Software
- ▶ VERITAS Volume Manager Version 3.2.0.0 Evaluation Software

Foundation Suite High Availability (HA) adds robust, easy-to-manage cluster administration to the disk and data management capacities already available in VERITAS Foundation Suite. The HA version includes:

- ▶ VERITAS Cluster Server Version 2.0.0.0, Evaluation Software

Additional product information can be found at:

<http://www.veritas.com/products/listing/ProductListing.jhtml>

AIX iSCSI Initiator Version 1.0

AIX iSCSI Initiator Version 1.0 provides the capability to send SCSI commands and receive responses over TCP/IP, enabling access to SCSI storage subsystems over TCP/IP/Ethernet. Additional product information may be found at:

<http://www.iscsistorage.com/>

<http://www.ietf.org/html.charters/ips-charter.html>

13.3.7 Contents summary

The following sections describe the contents or the additional media.

AIX 5L Expansion Pack

The AIX 5L Expansion Pack contains the following programs available worldwide. Some of these programs contain encryption as indicated.

The AIX 5L Expansion Packs contain the following programs shown in Table 13-17.

Table 13-17 AIX 5L Expansion Pack

Program name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
Certificate Authentication Services	N/A	Version 5.2.0.10
Cryptographic Library, 256-bit encryption	N/A	Version 5.2
Data Encryption Standard (DES) Library Routines for AIX, 64-bit encryption	Version 5.1.0.1	Version 5.2.0.1
IBM Directory Server and Client Utilities for Maximum Encryption (128-bit, Triple DES encryption)	N/A	Version 4.1.0.0
IBM Global Security Toolkit (Triple DES encryption)	Version 4.0.3.183	N/A

IBM Global Security Toolkit (Triple DES encryption)	Version 5.0.4.92	Version 5.0.4.92
IBM HTTP Server (128-bit encryption)	Version 1.3.19.4	Version 1.3.19.4
IBM IP Security (56-bit, Triple DES encryption)	Version 5.1	Version 5.2
Netscape Communicator (128-bit encryption)	Version 4.8	Version 4.8
Network Authentication Service (Triple DES encryption)	Version 1.3	Version 1.3
RSCT Encryption Support (56-bit DES, 56-bit Triple DES, 256-bit AES, 512-bit RSA, and 1024-bit RSA encryption)	N/A	Version 2.3
SecureWay Directory Server and Client Utilities for maximum encryption (128-bit, Triple DES encryption) includes a Software Developer Kit (SDK) that provides 128-bit encryption supporting Java Naming and Directory Interface (JNDI)	Version 3.2.2	N/A
SNMPv3 Encryption Support (CBC 56-bit DES encryption)	N/A	Version 5.2
Tools to Build Secure Java Applications (contains encryption)	Version 1.3.0	Version 1.3.1
IBM Web-based System Manager Security (128-bit encryption)	Version 5.1.0.50	Version 5.2.0.10

AIX 5L Bonus Pack

The AIX 5L Bonus Pack contains the following programs available worldwide. Some of these programs contain encryption as indicated.

The AIX 5L Bonus Pack contain the programs shown in Table 13-18.

Table 13-18 AIX 5L Bonus Pack

Product name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
Adobe Acrobat Reader with Asian fonts (Chinese and Korean)	5.06	5.06

Product name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
IBM AIX Developer Kit, Java2 Technology Edition, 32-bit version for POWER (contains encryption)	Version 1.4.1	Version 1.4.1
IBM AIX Developer Kit, Java2 Technology Edition, 64-bit version for POWER (contains encryption)	Version 1.4.1	Version 1.4.1
AIX Fast Connect, Evaluation Software	Version 3.1.2	Version 3.1.2
AIX iSCSI Initiator (see the following Web sites for more information): http://www.iscsistorage.com/ http://ietf.org/html.charters/ips-charter.html http://searchstorage.techtarget.com/whitepapers/	N/A	Version 1.0.0.0
Geodesic Systems Great Circle, Evaluation Software	Version 6.0.1.5	Version 6.0.1.5
IBM Directory Server	Version 5.1	Version 5.1
IBM Directory Server and Client Utilities for Maximum Encryption (128-bit, Triple DES encryption)	Version 5.1	Version 5.1
IBM Global Security Toolkit (Gskit) (Triple DES encryption)	Version 6.0.5.15	Version 6.0.5.15
Modular I/O Library	Version 2.3.0.0	Version 2.3.0.0
Open Secure Shell (OpenSSH) (Triple DES, Blowfish encryption)	Version 3.4.0.0	Version 3.4.0.52
IBM Tivoli Storage Resource Manager, Evaluation Software	Version 1.2	N/A
IBM Tivoli Storage Manager, Evaluation Software	Version 5.1	Version 5.1.5
VERITAS File System Evaluation Software	Version 3.4.2.2	N/A
VERITAS Volume Manager Evaluation Software	Version 3.2.0.0	N/A
VERITAS Cluster Server Evaluation Software	Version 2.0.0.0	N/A

Product name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
WebSphere Application Server, Trial Program	Version 5.01	Version 5.01
Wireless Application Protocol	N/A	Version 1.1.0.0

AIX 5L Web Download Pack

The AIX 5L Web Download Pack contains the following programs available worldwide. From this page, you will be able to select programs for download.

Table 13-19 AIX 5L Web Download Pack

Product name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
Adobe Acrobat Reader with Asian fonts (Chinese and Korean)	5.06	5.06
IBM AIX Developer Kit, Java 2 Technology Edition, 32-bit version for POWER (contains encryption)	Version 1.3.1	Version 1.3.1
IBM AIX Developer Kit, Java 2 Technology Edition, 64-bit version for POWER (contains encryption)	Version 1.3.1	Version 1.3.1
IBM 32-bit SDK for AIX, Java 2 Technology Edition, (contains encryption)	Version 1.4.1	Version 1.4.1
IBM 64-bit SDK for AIX, Java 2 Technology Edition, (contains encryption)	Version 1.4.1	Version 1.4.1
AIX Fast Connect Evaluation Software	Version 3.1.2	Version 3.1.2
AAIX iSCSI Initiator (see the following Web sites for more information): http://www.iscsistorage.com/ http://ietf.org/html.charters/ips-charter.html http://searchstorage.techtarget.com/whitepapers/	N/A	Version 1.0.0.0
Geodesic Systems Great Circle Evaluation Software	Version 6.0.1.5	Version 6.0.1.5
IBM Directory Server (128-bit, Triple DES encryption)	Version 4.1	Version 4.1

Product name	Operating systems	
	AIX Version 5.1	AIX Version 5.2
IBM Directory Server (128-bit, Triple DES encryption)	Version 5.1	Version 5.1
Mozilla for AIX Web Browser, 128-bit encryption	Version 1.4	Version 1.4
Netscape Communicator, 128-bit encryption	4.8	4.8
Netscape 7 for AIX, 128-bit encryption	7.0	7.0
Modular I/O Library	Version 2.3.0.0	Version 2.3.0.0
Open Secure Shell (OpenSSH) (Triple DES, Blowfish encryption)	Version 3.6.0.0	Version 3.6.0.52
IBM Tivoli Storage Manager, Evaluation Software	Version 5.1.5	Version 5.1.5
VERITAS File System Evaluation Software	Version 3.4.2.2	N/A
VERITAS Volume Manager Evaluation Software	Version 3.2.0.0	N/A
VERITAS Cluster Server Evaluation Software	Version 2.0.0.0	N/A
WebSphere Application Server, Evaluation Software	Version 5.02	Version 5.02
Wireless Application Protocol	N/A	Version 1.1

13.3.8 How to order

An Expansion Pack and a Bonus Pack are included at no additional charge with every new order of AIX 5L licenses when media is selected. Customers who already have the AIX 5L Expansion Pack or Bonus Pack may request a current Expansion Pack or Bonus Pack by contacting your IBM Representative, IBM Business Partner, or by shopping our online store. An additional media charge for Expansion Pack or Bonus Pack updates may be required in some countries. Programs on the Web Download Pack are always available for download, and updates to programs are not dependent upon availability of the AIX 5L Expansion Pack or Bonus Pack.

To receive the AIX 5L Expansion Pack or Bonus Pack, place an order using the information from Table 13-20.

Table 13-20 Product feature codes

Product	Feature Codes for 5692-A5L SPO for AIX 5L Version 5.1	Feature Codes for 5692-A5L SPO for AIX 5L Version 5.2
Expansion Pack includes some products with 128-bit or stronger encryption.	0921 Worldwide	0949 Worldwide
Bonus Pack includes some products with 128-bit or stronger encryption.	0920 Worldwide	0948 Worldwide

13.4 Linux

This section introduces the features and products of Linux and AIX affinity with Linux.

13.4.1 Introduction to Linux

The following sections detail IBM's commitment to Linux.

History of Linux

Linux is an operating system that is based on a development approach that delivers innovation and portability. Linux is an open, reliable, and efficient operating system that runs on virtually any platform from embedded systems to mainframes. Linux is the creation of Linus Torvalds, a Finnish computer science student, who developed it while he was a student at the University of Helsinki in 1991. The architecture is similar to the UNIX operating system. It provides a “free”, UNIX-like operating system solution across many computer architectures. After doing the initial development work, Torvalds made the source code available on the Internet for use, feedback and further development by others who were interested in helping to evolve Linux.

As an Open Source technology, Linux is not owned or controlled by any individual or company, but rather it is maintained by the Open Source community, which is a dedicated group of independent developers collaborating to make it the most open operating system. Being Open Source, the Linux operating system source code, like other Open Source technologies, can be acquired at no cost.

The GNU Project was launched in 1984 to develop a complete clone of the UNIX operating system that would be free software, that is, the GNU system. Variants of the GNU operating environment that use the Linux kernel are now widely used. For more information on the GNU Project, visit the Web site:

<http://www.gnu.org/gnu/the-gnu-project.html>

Customers are benefiting from the rapid innovation and enhancements made to Linux, enabled by the Open Source development approach. Linux is licensed under the terms of the GNU General Public License or GPL. The GPL requires, among other things, that the source code be made freely available to all who receive the program and that all modifications to the code be licensed using the GPL as well. This ensures that all changes and even derivative works remain Open Source. As a result, innovations are rapidly fed back into Linux for the benefit of all users. The following Web site provides more information on the GPL:

<http://www.fsf.org/copyleft/gpl.html>

While Linux is UNIX-like operating system, it is *not* the same as UNIX. The similarity begins and ends with the fact that Linux is based on the same design principles and standards as UNIX and it is derived from that heritage. The Linux source code is distinct from that of UNIX, and offers compatibility, portability, and horizontal scalability across all platforms.

Today, UNIX has split into a series of competing operating systems derived from the original code. Standards such as POSIX and UNIX 98 have been promulgated to specify many of the APIs and features of the various UNIX offerings. Linux is a single source operating system available to all. Through the GPL, developers must contribute their modifications back to the community.

IBM and Linux

IBM has made an expansive commitment to support Linux as an open computing environment. IBM understands that the open computing business model requires customer flexibility and choice. Linux is the epitome of both, at least in terms of operating systems. Linux continues to scale and address larger computing tasks, and IBM is doing its part to speed this process along, while optimizing IBM @server systems to offer customers the option of using Linux.

IBM provides continued support and participation throughout the world. You can learn more about this support and IBM's commitment to Linux at the following Web sites:

- ▶ The IBM Linux portal, which is a general point of entry into IBM and Linux:

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

- ▶ IBM Linux Technology Center (LTC):
<http://www.ibm.com/linux/ltc>
- ▶ IBM Solution Partnership Centers:
<http://www.developer.ibm.com/spc/index.html>
- ▶ IBM Linux support line:
http://www.ibm.com/services/e-business/linux_8.html
- ▶ The Open Source Development Lab:
<http://www.osdl.org/>

As part of this continuing commitment, IBM has teamed with leading commercial Linux distributors, including The SCO Group (previously known as Caldera Systems), Red Hat, SuSE, and TurboLinux. This cooperation allows IBM to port, test, and certify the performance of its offerings running on various Linux distributions. The result is to enable you to exploit the full potential of Linux.

Figure 13-2 shows IBM's strategy of the UNIX-based operating system environment. The direction for the IBM UNIX-based operation system evolution will be integration, coexistence, and interoperability of AIX and Linux, and a solid Linux affinity with the AIX 5L operating system.

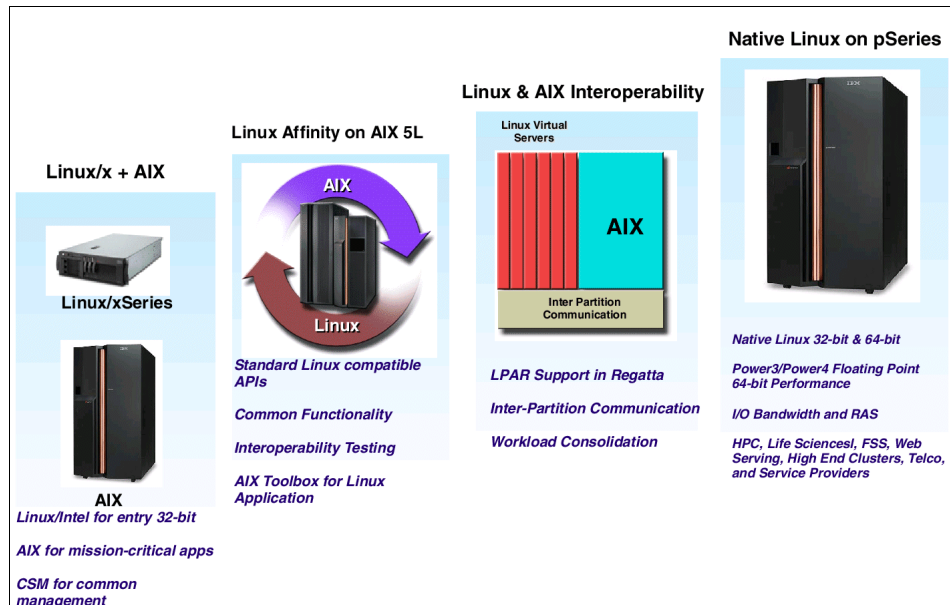


Figure 13-2 AIX and Linux with pSeries

Linux operating system and components

Over the past few years, the Linux operating system has become a real and viable alternative for PC users as well as corporate servers and users. Linux delivers the power and flexibility of a UNIX server or desktop. It also provides a set of utilities, Internet applications, and a fully functional desktop interface.

The Linux operating system has become a server platform for very powerful Internet applications. The Linux operating system is capable of running corporate Web servers, FTP, Gopher, and wide-area information server (WAIS) Web sites. Linux is a fully functional operating system similar to a UNIX system. It has all the standard features of enterprise UNIX systems. Management of the command structure is enabled through shells. We mention two of the many advanced shells that are available for Linux: the Bourne Again Shell (bash), and the tsch shell. Both shells support a complete shell-based programming language that you can use to create shell scripts.

There are four fundamental components of Linux:

Kernel	Provides program and hardware device management (such as disks and printers).
Shell	Receives commands from the user and sends those commands to the kernel for the execution through a standard user interface.
File structure	Files are organized into directories with the disk hardware. Each directory may contain any number of subdirectories each holding files. This is a similar structure to PC operating system file structures.
Utilities	Such as configuration and installation programs, firewalls, Web browsers, and Web server programs.

GNU project

By the 1980s, most operating systems were proprietary, which meant that you had to use the operating system provided for your specific hardware platform. The initiative of the Free Software Foundation (FSF) and the GNU Project motivated and stimulated open development and worldwide user cooperation. The main goal of the GNU Project was to develop a UNIX-compatible operating system named GNU (GNU is not UNIX), capable of running on various hardware architectures. Calling it GNU was a way of paying tribute to UNIX-like systems while saying that GNU was something different. It was to be 100 percent free, which meant that users would be free to redistribute the whole system, and free to change and contribute to any part of it. It was decided to make it UNIX-compatible because UNIX had already been proven in terms of design and portability.

The GNU Project was founded by Richard Stallman, the founder of the Free Software Foundation, author of the GNU General Public License, and the original developer of some GNU software programs, for example, the gcc compiler and the Emacs text editor. It took many years of hard work to write all the pieces of the GNU-based operating system, hundreds of programmers worldwide, and many hackers who worked very hard on the code and at the same time also used it.

The Linux kernel

By 1990, most of the software pieces had been written except for the most important one: the kernel. The kernel is the core of the operating system that loads first, and it remains in main memory. The kernel is responsible for memory management, process and task management, I/O, and disk management. It manages files and is the part of the operating system that interfaces with hardware devices. The open source portion of Linux that will always remain free is the kernel. Hundreds of developers submit their proposed changes to the kernel on an on-going basis. Linus Torvalds and his committee must approve these changes for them to be incorporated into the next kernel version. Today, the combination of GNU tools and commands and the Linux kernel is widely used around the world, and its popularity grows on a daily basis.

You can download the latest Linux kernel from:

<http://www.kernel.org>

Distributors

Linux is only the core of the operating system or the kernel. Linux distributions are prepackaged collections of software, generally comprising the kernel and various other packages.

As Linux gained popularity, a number of companies formed to distribute the Linux operating system along with a variety of additional value-added software packages and services. A distributor is a company who packages a Linux distributor for commercial purposes. They technically do not resell the software license, but rather sell a subscription service for support and maintenance. The major distributors include a graphical desktop environment and an installation program to help you get started. A Linux distribution also includes the Linux kernel plus utilities, programming tools, window managers, and other software that make up a full operating system. The kernel code in each distribution may be different, since distributions pull Linux versions. Linux Standards Base attempts to standardize the aspects of a distribution to minimize the changes that affect portability.

There are now over 100 companies doing various not-for-profit and for-profit distributions for a variety of hardware platforms. IBM has engaged Red Hat,

SuSE, and TurboLinux as Linux Distribution Partners (LDPs) to deliver the appropriate Linux solutions that support IBM's various hardware and software platforms.

The following sections describe the Linux distributors that are working with IBM to provide and support Linux for pSeries. Each distributor is wholly responsible for the contents, availability, and pricing of their offering. Regardless of how a Linux distribution is ordered, the distributor is responsible for maintenance and support.

Red Hat Linux

Red Hat is the best known of the Linux distributors. Red Hat 7.1 was announced for pSeries in November 2001 and made generally available in March 2002. An update supporting more pSeries systems became generally available in December 2002 and is designated as "Red Hat Linux 7.1 for pSeries (64-bit)". It is a 64-bit kernel and supports 32-bit applications. Details on the offering, including pricing, can be found at:

<http://www.redhat.com/software/eserver/pseries>

IBM has recently signed an agreement with Red Hat that aims to bring the Red Hat Enterprise Linux Advanced Server product to pSeries in 3Q2003. This will be a full 64-bit kernel with 32-bit and 64-bit application support.

SuSE Linux

SuSE was the first of the IBM Linux Distribution Partners to release Linux for the pSeries and RS/6000. SuSE Linux for PowerPC Version 6.4 was released in June 2000. It supported the RS/6000 B50, 150, and F50 systems. Version 6.4 was based on the 32-bit Linux 2.2.x kernel. It was superseded by Version 7.x, which is based on the 32-bit Linux 2.4 kernel. The latest version of SuSE for enterprise customers, SuSE Linux Enterprise Server 8 (SLES 8), is now available for pSeries. It became available in December 2002 and contains the 64-bit Linux kernel (Version 2.4.19) as well as both 32-bit and 64-bit application support. SLES 8 is the first distribution to be branded as "Powered by UnitedLinux Version 1.0" (details on UnitedLinux are provided below).

Full details on SLES 8 for pSeries are available directly from SuSE at:

http://www.suse.com/us/business/products/server/sles/i_pseries.html

United Linux

In June 2002, four of the Linux distributors formed the UnitedLinux organization. The UnitedLinux companies are Conectiva S.A., The SCO Group, SuSE Linux AG, and TurboLinux, Inc. They have pooled their development resources to create and distribute a common, standards-based server distribution of Linux

focused on the business customer. The distribution supports Intel processor-based systems and each distributor decides on which other platforms, including IBM @server pSeries systems, to support.

While the development organization is common, each member will sell and service its own distribution at whatever price it chooses. More information can be found at their Web site at:

<http://www.unitedlinux.com>

SuSE Linux Enterprise Server 8 for pSeries carries the UnitedLinux “Powered by UnitedLinux Version 1.0” brand. Since the UnitedLinux brand is intended to ensure that the distributions are common across all UnitedLinux branded offerings, each of the other distributions will be equivalent to SLES 8 at the programming interface level.

TurboLinux

TurboLinux has a strong presence in the Asia Pacific market. TurboLinux 6.5 for pSeries is currently available and supports the p640-B80 with a 32-bit kernel. TurboLinux plans to provide an updated pSeries distribution based on UnitedLinux Version 1.0 during 2003.

For more information, visit the Web site:

<http://www.turbolinux.com/products/eserver/#pSeries>

13.4.2 Linux and pSeries

Linux for pSeries is a key element of the IBM @server Linux strategy. During the past few years, IBM has been developing the ability to integrate Linux capabilities into the pSeries systems. pSeries systems can run Linux alone or can run both IBM AIX and Linux in logical partitions (LPAR) of the same server. In this section, we will discuss Linux enablement for the pSeries and the ways that you can integrate Linux applications into commercial pSeries environments.

Linux supported pSeries systems

Both 32-bit and 64-bit versions of Linux for pSeries are being provided to optimize customer choices and exploit pSeries hardware capabilities. The table below details the Linux distribution and status LPAR support for pSeries 64-bit systems. SuSE Linux Enterprise Server 8 has a 64-bit kernel with support for either 32-bit or 64-bit applications. Red Hat Linux 7.1 for pSeries (64 bit) provides a 64-bit kernel, but the tools and libraries support only 32-bit applications. In practice, this is not a major limitation, as the majority of today’s Linux applications are 32-bit. Database managers and many high performance

computing applications tend to make use of 64-bit. You may refer to the Table 13-21 for summarized information on Linux supported pSeries systems.

Table 13-21 Linux distribution support for pSeries 64-bit systems

Models	SuSE Linux Enterprise Server 8 (UnitedLinux Version 1.0)	Red Hat Linux 7.1 for pSeries (64-bit)	Linux static LPAR support
170 (7044-170)	Yes	No	No
270 (7044-270)	Yes	No	No
p610 (7028-6C1) p610 (7028-6E1)	Yes	No	No
p615 (7029-6C3) p615 (7029-6E3)	Yes	No	No
p620 (7025-6F0) p620 (7025-6F1)	Yes	Yes	No
p630 (7028-6C4) P630 (7028-6E4)	Yes	No	Yes with SLES 8
P640 (7026-B80)	Yes	Yes	No
P650 (7038-6M2)	Yes	No	Yes with SLES 8
P655 (7039-651)	-	No	Yes with SLES 8
P660 (7026-6H0) P660 (7026-6H1)	Yes	Yes	No
P660 (7026-6M1)	Yes	Yes	No
P670 (7040-671)	Yes	No	Yes with SLES 8
P690 (7040-681)	Yes	No	Yes with SLES 8

Note:

1) **For p615:** A distribution of SuSE is available at the time this redbook was written. In Japan, TurboLinux is also available.

2) **For p630 and p650:** These LPAR-capable systems are supported either with or without the use of LPAR. Only statically-configured LPARs are supported, meaning that Linux must be stopped and restarted in order to change the partition configuration. These systems may be ordered as Linux ready Express Configurations at special prices. Orders for SLES 8 may be placed through IBM at the time of purchase for Linux ready Express Configurations.

3) **For p670 and p690:** These LPAR-capable systems are supported only in LPAR mode. Only statically-configured LPARs are supported, meaning that Linux must be stopped and restarted in order to change the partition configuration. Because Linux currently does not scale well beyond 6- to 8-way SMP, Linux is only supported in an LPAR on these large systems. A maximum of eight processors is recommended for each Linux LPAR on these systems. While Linux can be run successfully on systems/LPARs with more than eight processors, typical application workloads will only effectively utilize the equivalent of four to eight processors.

4) **For p655:** Linux on the p655 is supported only on systems running in LPAR mode in conjunction with AIX in a second LPAR. Support for Linux running without LPAR (native) and in one or more LPARs without AIX will be available in 3Q 2003.

I/O device and adapter support

There are a large number of adapters and devices that can be attached to pSeries and RS/6000 systems running AIX. While some of the devices, for example, PCI adapters, have Linux drivers for Intel as well as AIX drivers, these cannot be utilized directly in Linux for pSeries.

As part of the base enablement of Linux for pSeries, adapters are enabled in addition to the base device support like SCSI, Ethernet, and so on. The supported adapters are detailed in the Linux for IBM @server pSeries Facts and Features document available at:

http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf

Other adapters not currently supported will obviously be required and some are already ported but not yet tested. These will be included in future Linux for pSeries distributions and work is being done to help provide others as part of special bids to customers interested in testing/deploying Linux for pSeries.

Configuration examples

The pSeries offers outstanding server performance and a great server consolidation platform. These servers provide the ability to run AIX or Linux as single operating systems or multiple LPARs running any combination of AIX and Linux. Figure 13-3 shows an overview of the internals of the AIX and Linux operating systems, and how applications can run on pSeries in three different scenarios.

Scenario 1 shows the traditional AIX-only-based pSeries environment. All AIX versions prior to Version 4.3.3 or 5L only run applications compiled for AIX. That means Linux applications cannot be run, unless they are explicitly ported to AIX.

Scenario 2 shows how AIX Version 4.3.3 and AIX 5L incorporate Linux in the pSeries by means of the AIX toolbox for Linux applications. The AIX toolbox for Linux applications provides the capability to recompile open source applications to run as AIX applications. In most cases, it does this without porting and using the same open source build environment as Linux. The AIX toolbox for Linux applications are discussed in 13.4.3, “Linux affinity on AIX 5L” on page 602.

Scenario 3 shows the fully Linux-enabled scheme. It has a Linux for pSeries kernel. Customers can integrate a broad spectrum of Linux applications into their pSeries servers. The customer can generate Linux code for any Linux application available for the pSeries.

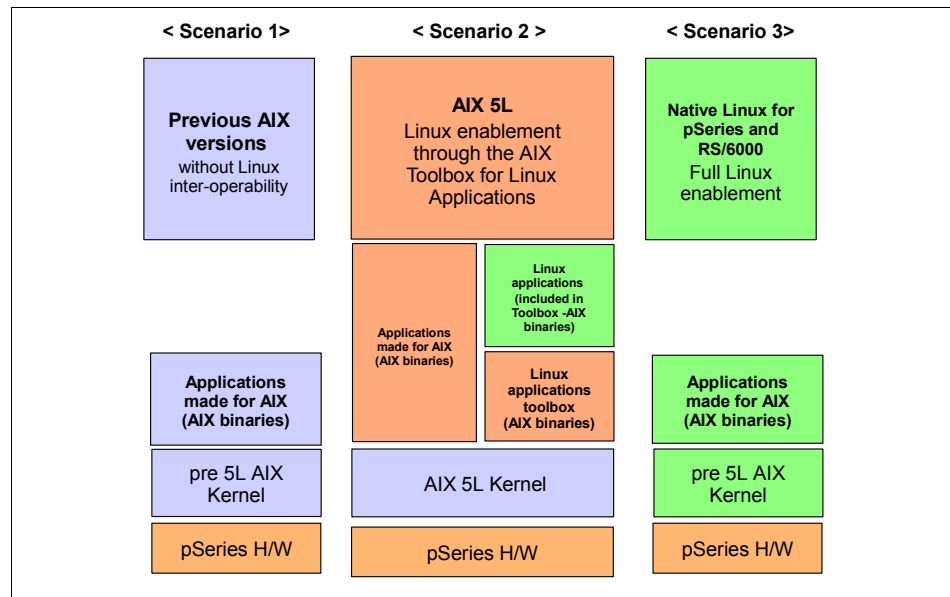


Figure 13-3 Configuration example

Linux for pSeries scalability

Linux kernel 2.4 has proven to scale well to up to eight processors in a symmetric multiprocessing (SMP) system depending on the workload. This makes it a good match for previous RS/6000 systems (such as the 170 and 270), pSeries systems (such as the 610, 615, 620, 630, 640, 650, 655, and 660), and a 1- to 8-way LPARs on the 670 and 690. Linux for pSeries can also scale up to 64 GB in memory and 2 TB in file system size.

Linux for pSeries RAS

A key attribute of the pSeries is its critical Reliability, Availability, and Serviceability (RAS) features. Drawing from IBM's autonomic computing efforts, pSeries continues to enhance the scope of its RAS capabilities. However, implementation may vary, depending on the platform and environment on which you are working. For more information about RAS for the pSeries, go to:

http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_ras.html

Linux for pSeries is rapidly developing RAS capabilities. However, many pSeries RAS features are only fully realized when running AIX. The following pSeries RAS features are supported when running Linux:

- ▶ Chipkill and ECC memory
- ▶ Disk mirroring (Software)
- ▶ Journaled file system (several available under Linux)
- ▶ PCI Extended Error detection
- ▶ Redundant, hot-plug power, and cooling
- ▶ Error logging
- ▶ Service processor (partial)
- ▶ Boot-time processor and memory deallocation
- ▶ First Failure Data Capture (except for I/O)

Some of the pSeries RAS features that are not currently supported in Linux include:

- ▶ Hot-swapping of disk drives
- ▶ Hot-plug PCI and memory
- ▶ Dynamic Processor Deallocation
- ▶ PCI Extended Error recovery
- ▶ Error reporting to Service Focal Point
- ▶ Error log analysis

- ▶ Remote support
- ▶ High Availability Cluster Multiprocessing (HACMP) (Alternative third-party and open source solutions may be available.)

Logical partitioning

The LPAR capabilities of the pSeries server make it possible to run one or more instances of Linux along with AIX. This provides a convenient way to begin developing and deploying Linux operating system-ready applications as desired, while retaining the enterprise-ready capabilities of AIX for mission-critical or highly-scalable workloads. LPAR also allows large pSeries servers to be partitioned to run Linux only workloads.

Let us review some key terminology:

- ▶ LPAR: This is the shortened version of the familiar term logical partitions. It is widely used to indicate that a server can run multiple separate operating system instances.
- ▶ Static LPAR: This term applies to a pSeries server or partition. It means Linux must be stopped and restarted in order to change the partition configuration. All pSeries Linux partitions are static.
- ▶ Dynamic LPAR (DLPAR): This term indicates that a partition accepts certain rearrangements of resources. Partitions running AIX and AIX 5L applications can have dynamic movement of processors and memory between partitions. DLPAR is not currently supported by Linux; however, Linux partitions can be created on systems enabled for dynamic LPAR.
- ▶ Hypervisor: This is the underlying control mechanism that sits between the physical hardware and the operating systems. The hypervisor owns all system resources and creates partitions by allocating these resources and sharing them.

AIX is not required to partition the machine, and it is not required to install and run Linux. Nevertheless, AIX provides much greater features for hardware analysis, intermittent fault detection, and error logging.

For more information, visit the pSeries Web site at:

<http://www.ibm.com/servers/eserver/pseries/linux/>

AIX and Linux can run in separate partitions at the same time on an LPAR-enabled pSeries server in any combination, that is, zero or more Linux partitions along with zero or more AIX partitions. This allows you to consolidate workloads from several separate servers onto a single system.

Since partitioning in an LPAR-supported pSeries is controlled by the firmware-based hypervisor and the Hardware Management Console (HMC), AIX is not required to run Linux.

Figure 12-3 shows a typical logically partitioned pSeries server. It shows the major components the HMC linked to both the partitions and the hypervisor. This includes the internal components that the HMC can control, processors, memory, and I/O adapters. The logical partitioning functions allow the rearrangement of resources to and from a partition, as well as to move resources between two partitions, without operating system reboot. The new Dynamic LPAR function in pSeries can dynamically move resources for AIX partitions. It is currently not supported by Linux. However, Linux partitions and dynamic LPAR can run on the same server.

For more information about LPAR-supported pSeries servers, see:

<http://www.ibm.com/servers/eserver/pseries/lpar/wp.html>

In the example in Figure 13-4, any Linux partition can be an AIX partition.

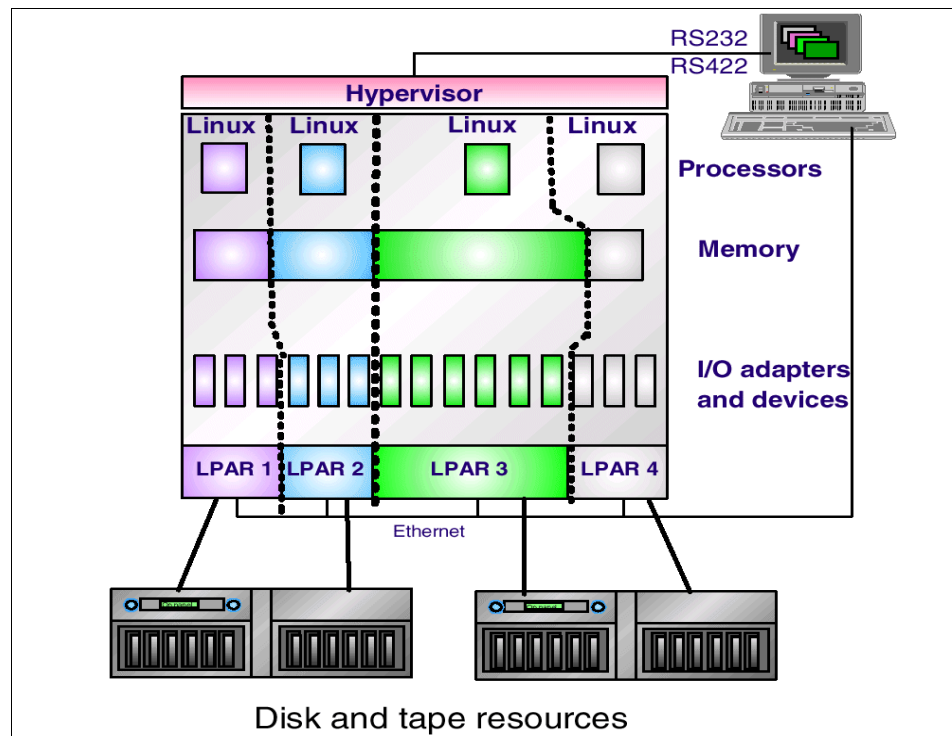


Figure 13-4 Linux partitions on the pSeries system

Linux for pSeries clusters

Currently, none of the IBM software that has been announced for IBM @server Cluster products, the IBM @server Cluster 1350 and Cluster 1600, has been ported to Linux on pSeries. IBM has announced a Statement of Direction to port its Cluster Systems Management (CSM) technology and other tools to Linux for pSeries.

The Beowulf clustering technology and other Open Source and some commercial products can be used to cluster pSeries systems running Linux to provide compute or high-availability clusters. Beowulf has its Web site at:

<http://www.beowulf.org>

Myricom has the Myrinet switch available for Linux for pSeries. It can be used as a high-speed interconnect to cluster systems of pSeries machines running Linux. Gigabit or 10/100 Ethernet connections can also be used. For more information on Myricom, refer to the following Web site:

<http://www.myricom.com>

Linux ready Express Configurations

To make it easy to get started with Linux for pSeries, IBM has introduced a number of pSeries Linux ready Express Configurations. These systems represent some of the most popular configurations of the p630 and p650 systems. They are provided without an AIX license and offer great savings with the ability to add additional features. Linux can be ordered through IBM at the time of initial purchase of these systems.

Check the Linux for pSeries Web site for current configurations and pricing at:

<http://www.ibm.com/servers/eserver/pseries/hardware/express.html>

13.4.3 Linux affinity on AIX 5L

In this section, we discuss the Linux affinity that AIX 5L supports.

Overview

IBM's plans for AIX affinity with Linux are implemented in two phases:

- ▶ The first phase is the release of the AIX toolbox for Linux applications, which is available at:

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

The AIX toolbox for Linux applications contains GNU and other commonly used tools helpful for recompiling an application for use on AIX.

- ▶ The second phase is IBM's planned inclusion of additional Linux compatible APIs and header files in AIX 5L Version 5.1 and 5.2.

AIX affinity with Linux uses a Application Programming Interface (API) approach to providing Linux application interoperability. This approach is *not* an environment or an additional layer or wrapper to run Linux applications in or on. It is the integration of Linux compatible APIs and header files into AIX 5L. Therefore, recompiled Linux applications are treated as native AIX applications and have access to all the reliability, scalability and availability of AIX. The result is a tighter integration of the application to the operating system than can be achieved with an Application Binary Interface (ABI) approach.

AIX has been developed using UNIX industry standards and, as such, there is a high degree of compatibility at the API level between AIX and Linux. This degree of similarity is such that many Linux applications can be recompiled and run on AIX Version 4.3.3 and AIX 5L Version 5.0 today using the AIX toolbox for Linux applications. AIX 5L Version 5.1 and 5.2 have even more powerful interoperability between AIX and Linux by including APIs that are presently not similar between AIX and Linux, resulting in an even higher degree of Linux application compatibility. Thus, much of the functionality of AIX 5L Version 5.1 and 5.2 already is available in AIX Version 4.3.3 and AIX 5L Version 5.0, with full implementation of IBM's Linux affinity available in AIX 5L Version 5.1 and 5.2.

Performance considerations

When it comes to the best utilization of AIX affinity with Linux, it is important to consider impacts to performance. AIX affinity with Linux is designed to provide the best performance possible; however, there are a couple of issues to consider that are outside the control of AIX affinity with Linux that can influence performance.

The Linux application being deployed on AIX will have full access to all AIX functionality, just like an application natively developed for AIX. AIX has a high level of compatibility with Linux, and IBM is providing an even more powerful affinity between AIX and Linux with AIX 5L Version 5.2. Therefore, for a Linux application to take advantage of AIX, it does not need to run through any additional layer or wrapper. The question of performance is not one of the functionality of the recompiled Linux application to take advantage of AIX and the IBM POWER architecture, but one of the performance of the compiler used to build the application. Most applications that have been developed natively for AIX use the IBM Visual Age compiler, while applications developed natively for Linux utilize the GNU compilers. Thus, you can expect to see a performance advantage for AIX applications that have been built using the IBM Visual Age compiler. Unfortunately, however, the IBM Visual Age compiler is not available for Linux applications at this time.

The Application Programming Interface (API) method that AIX utilizes provides a higher degree of integration between the application and the operating system than can be achieved using a layered or wrapper approach, such as found in an Application Binary Interface (ABI) approach.

AIX toolbox for Linux applications

AIX toolbox for Linux applications is a group of GNU and open source tools and utilities for building and deploying Linux applications on AIX. It includes a collection of programs that have already been recompiled and tested.

The goal of the AIX toolbox for Linux applications is to provide ready-to-run, installable Open Source tools and facilitate recompilation of Open Source Software, without modifications, on AIX systems. These days, many Open Source applications are created on Linux systems and are using Linux libraries and APIs. Since Linux is a UNIX clone and not a UNIX branded operating system, which means it aims towards POSIX and Single UNIX Specification compliance, supporting Linux Open Source Software requires that the Linux APIs be available in AIX to recompile the sources. For more information regarding the Linux kernel, please refer to the following Web site:

<http://www.kernel.org>

Once recompiled, the original Linux source applications become native AIX applications, which means they can take advantage of the same scalability, reliability, and performance as any other AIX application. These applications are AIX binaries.

The uses of AIX toolbox for Linux applications include:

- ▶ Build and package Linux applications for use on AIX
- ▶ Run Gnome and KDE desktops
- ▶ Run other popular software commonly found in Linux distributions
- ▶ Manage open source software using the popular RPM package management system
- ▶ Develop new applications for AIX using GNU and Linux application development tools

Figure 13-5 on page 605 shows the AIX toolbox structure running over the AIX kernel.

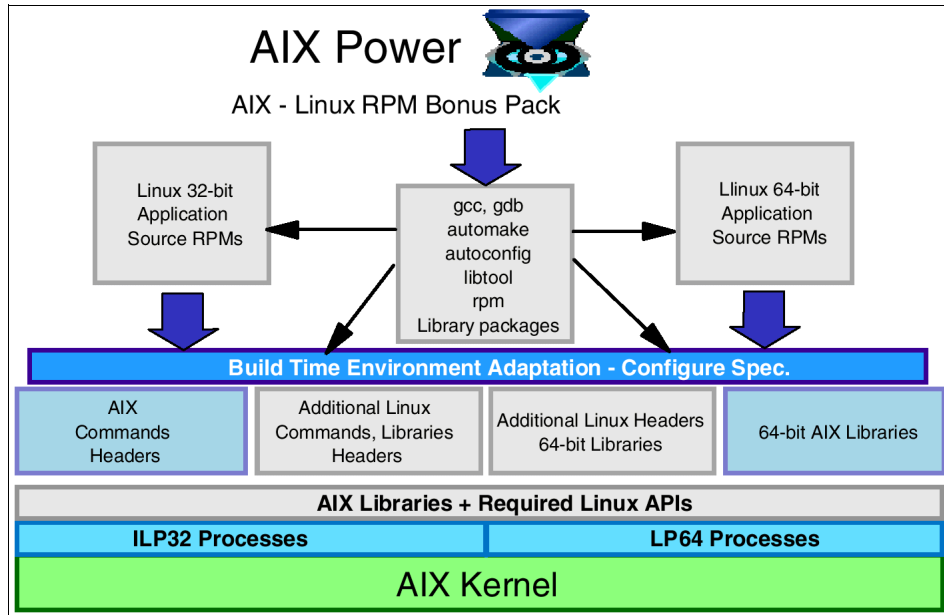


Figure 13-5 AIX toolbox for Linux applications

Open Source Software in the AIX Toolbox

The AIX toolbox for Linux applications provides the capability to easily compile and run Linux applications, providing the Linux applications full access to AIX functionality and features. The AIX toolbox for Linux applications contains Open Source applications, which includes recompiled versions and source code of the Gnome and KDE desktop environments, as well as systems utilities, including GNU application development tools, libraries, shells, utilities, and Graphical User Interface (GUI) desktops. The compiled Linux applications then become AIX binaries.

Here are some of the tools included in the AIX toolbox for Linux applications:

Graphical desktops KDE and Gnome desktops

Development tools automake, gcc, g++, gdb, make, autoconf, libtools, m4, patch, rpm-build, and more

Development libraries

libxml, db, gtk+, libtiff, libjpeg, ncurses, qt, zlib, and more

User interface for desktops

Sawfish and WindowMaker (window manager for X Window System), Koffice (a set of office application for

KDE), the Gnome control-center, enlightenment window manager, and more

Database application

MySQL, including client and libraries

System applications

Samba client, Gnome and KDE utility programs, GnoRPM, mtools, and more

System shells

bash, tcsh, zsh, mc, and the GNU utilities for shell scripts

Programming languages

Python, C and C++ compilers, and PHP

Internet applications

Fetchmail, lynx, ncftp, pine, xchat, ytalk, wget, elm, and more

Archiving applications

tar, cpio, cdrecord, zip, unzip, bzip2, gzip, and more

Graphical games

GNU chess program, KDE games, and the X Window graphical chessboard

Editing applications emacs, gnotepad+, vim, and more

The AIX toolbox for Linux applications media is included in all AIX media purchases. All the tools are packaged in RPM format. Aside from the Toolbox media that is shipped with AIX, you can also get a complete listing, or download all available tools included in the AIX toolbox for Linux applications, at the following URL:

<http://www.ibm.com/servers/aix/products/aixos/linux/download.html>

or directly from the FTP site:

<ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox>

The FTP site contains the new and old versions of the tools, the spec files, and the patches that were used to build the tools.

The AIX affinity for Linux provides easy migration of Linux applications to AIX systems using the GNUPro application development tools. With these tools, you can recompile Linux-based Open Source Software (OSS) and GNU software for use on AIX. The Toolbox provides all the tools necessary for effective software compiling and debugging.

Other sources of Open Source Software

Apart from the Toolbox, there are other sources where you can obtain Open Source Software for AIX on the Internet. Some distribute ready-to-run

executables of Open Source Software packages in Backup File Format (BFF), others in TAR format; normally, they are given a file name extension of .bff and .tar, respectively. The BFF format is the standard installation packaging format for AIX. Additional information about the individual files delivered in the product are included inside the BFF package. This packaging is named Licensed Product Package (LPP).

You can download a wide range of precompiled and packaged software for AIX from Open Source Software Web sites such as Group Bull at:

<http://www.bullfreeware.com>

or UCLA at:

<http://aixpdslib.seas.ucla.edu/aixpdslib.html>

13.4.4 Software service and support

Linux support is readily available from many sources. It ranges from free support from the Open Source community at large, to fee based service contracts with service organizations and Linux distributors, such as Red Hat and SuSE. Details on these offerings are available at the respective distributor's Web sites. Maintenance contracts for software upgrades can also be obtained from the distributors. Initial installation and maintenance is usually bundled into a Linux distributor's product.



Cluster, resource, and performance management software

This chapter introduces the following cluster software products dedicated to build, manage, and expand clusters efficiently:

- ▶ IBM Cluster Systems Management for AIX 5L (CSM)
- ▶ IBM Parallel System Support Programs for AIX 5L (PSSP)
- ▶ IBM General Parallel File System for AIX 5L (GPFS)

It also provides an overview of the resource management feature of AIX named Workload Manager for managing system resources.

The following performance management software products are also introduced:

- ▶ AIX Performance ToolBox for POWER and AIX Performance AIDE for POWER
- ▶ AIX Fast Connect for POWER

14.1 IBM Cluster Systems Management for AIX (CSM)

Cluster-ready software from IBM enables any multiple IBM @server system solution or cluster to look and act like a single system for end users and system administrators.

Cluster Systems Management (CSM) Version 1.3 is designed for simple, low-cost management of distributed and clustered IBM @server pSeries and xSeries servers. For organizations with both Linux and AIX applications, a single AIX 5L Version 5.2 management console can provide management services to AIX 5L and Linux clients in distributed and clustered configurations. CSM is also a key element of the IBM @server Cluster 1600 and Cluster 1350, platforms ideal for workload consolidation or for achieving high degrees of scalability and performance for applications that take advantage of clustered systems architectures, for example, computational modeling in high-performance computing or multi-terabyte data warehouses in large corporations.

CSM for AIX 5L delivers extensive clustering capabilities on Cluster 1600 platforms and clustered or distributed pSeries servers. Using a WebSM interface, a CSM management server on AIX 5L Version 5.2 can manage up to 128 operating system images. Scaling greater than 128 operating system images is available by special bid.

Workloads that benefit from clustering include:

- ▶ Large-scale computational models (public, industrial, and financial sectors)
- ▶ Large-scale databases (computer systems integrators, communications, distribution, and financial sectors)
- ▶ Consolidated UNIX and Linux workloads in large and medium data centers (computer systems integrators, public, industrial, financial, communications, distribution, and communications sectors)

14.1.1 Product positioning

Cluster Systems Management (CSM) Version 1.3 for AIX (refer to Software Announcement 202-262, dated October 8, 2002) should be marketed to new or existing cluster customers interested in consolidating AIX and Linux workloads or in remote management of operating system images installed on servers that are distributed across geographically dispersed data centers. It is not possible to mix PSSP and CSM in the same cluster.

14.1.2 CSM for AIX 5L tasks

CSM for AIX 5L performs the following tasks:

- ▶ Operating system and CSM installation on managed nodes

CSM can be configured to perform full operating system and CSM installation on all managed nodes, using installation management services provided by the operating system. For example, CSM for AIX interacts with the Network Installation Manager (NIM), and CSM for Red Hat Linux uses Kickstart. Full operating system installation requires remote hardware control support.

- ▶ Managing node and node-group information

CSM is designed to allow administrators to perform operations on both individual nodes as well as groups of nodes within a cluster. Administrators can create both static node groups, containing specifically identified nodes, as well as dynamic node groups, whose membership changes if attributes of nodes change.

- ▶ Remote hardware control

The CSM Management Server can be used to power on, power off, or query nodes. Additional functions for CSM for AIX 5L includes collecting adapter information, listing of nodes managed by an HMC, and starting a network boot for a node. Additional functions for CSM for Linux includes resetting either the Management Processor (MP) on the node or the Management Processor Adapter (MPA), listing the MPs that are managed by an MPA, configuring and enabling SNMP alerts supported by the MP, downloading and clearing the MP logs, and querying the node environmental information maintained by the MP. Remote console support is designed to allow you to open one or more consoles to a node serial port. This will allow the user to monitor installation or access a node's console directly.

- ▶ Running commands remotely

The distributed shell (**dsh**) command runs commands remotely across multiple nodes in parallel. Optionally, it can use an underlying remote shell that is specified by the user (for example, a remote secure shell that complies with the Internet Engineering Task Force (IETF) Secure Shell protocol). By default, **rsh** is used on AIX and **openssh** on Linux.

- ▶ Configuration File Manager (CFM)

CFM is designed to enable the administrator to distribute configuration files across the entire CSM cluster, and allows for variations in the files based on node groups and host name. All CFM configuration file distribution is initiated from the CSM Management Server and can be completely automated.

- ▶ Monitoring system events

A flexible distributed system monitoring application is provided by CSM. This monitoring application allows the administrator to define conditions of interest to monitor on the system. When a monitored condition arises, automated actions can take place, such as logging, e-mail, or paging.

A set of predefined conditions and responses is also provided with CSM, to monitor standard system resources such as file systems, programs, CPU and memory utilization, network availability, power status, and others.

- ▶ Security

Host security is provided by the operating system security controls (for example, file group/owner attributes and permission bits, file system ACLs, and effective and real user IDs). Many of the host's functions or modifications can only be performed by a root user. Remote shells that conform to the IETF Secure Shell protocol are designed so that they can be specified by using the **dsh** command, thus providing a greater level of security flexibility and control to the administrator when issuing commands to remote hosts.

- ▶ Distributed Command Execution Manager (DCEM)

DCEM gives you a GUI that is designed to let you run a command or script on multiple distributed machines on a network at the same time. You can specify a collection of individual machines, or you can create groups of machines and save them to use again. DCEM provides real-time command execution status on the individual machines that you specify, showing them in a waiting, working, successful, or failed state. It helps you to create, save, and edit command specifications, and it also creates a log of all distributed command activity.

14.1.3 CSM for AIX Version 1.3

CSM for AIX Version 1.3 provides the following features:

- ▶ CSM is designed to automatically set up security in the CSM cluster between the management server and the managed nodes, taking a great deal of manual steps away from the administrator. It provides the configuration of the cluster security services as noted above as well as automatic setup of rsh or openssh security.
- ▶ A GUI to CSM on AIX is provided by new applications/plugins that have been integrated into the Web-based System Manager (WebSM) Console.
- ▶ Software Maintenance System (SMS) is a collection of tools that are designed to enable an administrator to update, install, query, and remove Linux RPMs across an entire CSM Linux cluster. The SMS provides a simple way to update RPMs across multiple Linux machines with different RPM

configurations. SMS works across the entire cluster, on individual nodes, or on node groups.

- ▶ All of the CSM functionality with the exception of full operating system install is available for CSM Version 1.3 for Linux nodes added to a cluster managed by an AIX management server, which must be running AIX 5L Version 5.2.
- ▶ The administrator can run diagnostic probes provided by CSM to automatically perform health checks of particular software functions if a problem is suspected.

14.1.4 Hardware requirements

The hardware requirements are:

- ▶ Runs on IBM *e*server pSeries and RS/6000 servers and nodes with AIX 5L Version 5.1 or Version 5.2.
- ▶ A Hardware Management Console (HMC) support is required for remote hardware control. The HMC must be a Hardware Management Console for pSeries Version 1.3.1, or later. CSM remote hardware control is supported as follows:
 - Maximum of eight 690/670 systems per HMC
 - Maximum of sixteen p650, p655, or p630 systems per HMC
 - In mixed environments, maximum of 16 of any p650/p630/p655 configurations (for example, four p655, eight p650, and four p630)
 - Maximum of eight 690/670 and any eight p650/p630/p655 combinations
 - Maximum of 32 LPARs per HMC
- ▶ CSM provides distributed management services (but not remote hardware control or network operating system installs (OSIs)) for any machines running with AIX 5L.
- ▶ The following additional hardware and hardware specifications are required for CSM for AIX 5L Version 1.3:
 - On the CSM management server, one Ethernet adapter for each management virtual LAN (VLAN) used for hardware control, and one Ethernet adapter for each cluster VLAN used for CSM installation and administration
 - On each managed node, one Ethernet adapter for the cluster VLAN used for CSM installation and administration
 - On the management server, a minimum of 128 MB of memory and 120 MB of disk space to install CSM

- On each managed node, a minimum of 128 MB of memory and 20 MB of disk space to install and run CSM
- On the management server, 2 GB of disk space for each version of AIX 5L installation images
- On the management server and on each node, additional disk space for the base AIX 5L Version 5.2 operating system and filesets, or AIX 5L for POWER Version 5.1 with the 5100-03 Recommended Maintenance package.

14.1.5 Software requirements

The CSM single-point-of-control server is known as the management server and must be running AIX 5L Version 5.2 with Recommended Maintenance package 5200-01 (APAR IY39795). The other machines within the cluster are referred to as managed nodes and can be running AIX 5L Version 5.2 with 5200-01, or Version 5.1 with the 5100-04 Recommended Maintenance package (APAR IY39794).

For all AIX servers, the CSM for AIX 5L Version 1.3.1 services provided in Table 14-1 is required.

Table 14-1 CSM for AIX 5L Version 1.3.1 service required

Description	APAR number
Added Service for RSCT on AIX 5L Version 5.1	IY42782
Added Service for RSCT on AIX 5L Version 5.2	IY42783
CSM for AIX 5L added service	IY42353
CSM Support for p670 & p690 POWER4+	IY42356
CSM Support for 7026 servers, 9076 SP nodes and p650	IY42379
CSM Version 1.3.1 support for 9076 SP Node Features 2054/2058	IY42847
CSM Support of p655 Power4+	IY42377

Certain CSM for AIX 5L functions require non-IBM software. The following non-IBM software is required and can be obtained from the following sources:

- ▶ AutoUpdate Version 4.3.4 and Perl-libnet Version 1.0703 is necessary if you want to perform the software maintenance installation and upgrade of

non-CSM RPMs on your Linux managed nodes from the AIX management server. Available from the following Web site:

<http://freshmeat.net/projects/autoupdate>

- ▶ openCIMOM Version 7 is necessary if you want to perform Remote Hardware Control operations for IBM @server pSeries servers attached with a Hardware Management Console (HMC). You can get this software from the *AIX toolbox for Linux applications* CD (refer to Software Announcement 201-090, dated April 17, 2001), or download it from:

<http://www.ibm.com/servers/aix/products/aixos/linux/download.html>

14.2 IBM Parallel System Support Programs (PSSP)

PSSP software provides a comprehensive suite of applications for the installation, operation, management, and administration of the RS/6000 SP and IBM @server Cluster 1600 from a single point of control.

The PSSP suite of applications aids the operator/administrator in performing the tasks of installation, configuration management, operation, file management, user management, job management, accounting, change management, security, availability, problem management, and support of the high-speed interconnect fabric.

For many of the tasks listed above, PSSP enables the SP/cluster operator/administrator to avoid the problems of working with individual SP nodes and servers in a cluster or Cluster 1600 and their individual operating systems and applications. Instead, these cluster building blocks can be managed as a single system. Where the nodes and servers must be treated individually, PSSP tools reduce the amount of manual work that would otherwise be necessary.

Benefits: PSSP greatly improves operator/administrator productivity by enabling them, for many of the above tasks, to work with a single SP system or single cluster instead of a possibly large number of individual nodes and/or servers, and their individual operating systems and software.

14.2.1 Product positioning

Parallel System Support Programs (PSSP) VS3.5 should be marketed to all SP and IBM @server Cluster 1600 customers interested in taking advantage of its new or expanded features (for example, full support of the AIX 5L 64-bit kernel), while still retaining legacy cluster hardware that is supported only in PSSP-managed clusters.

Included among the servers that are supported in PSSP-managed clusters are:

- ▶ 332 MHz SMP SP thin and wide nodes, 200/375/450 MHz POWER3 SP thin and wide nodes
- ▶ 222 and 375 MHz POWER3 SP high nodes
- ▶ RS/6000 H80, M80, S70, S70 Advanced, and S80 servers
- ▶ IBM @server pSeries 630, 660, 670, 680, and 690 servers

In addition, the SP Switch and the SP Switch2 are supported only on PSSP-managed clusters.

14.2.2 PSSP Version 3.5 functional enhancements

The following enhancements are related to PSSP Version 3.5.

- ▶ Support is added for the following hardware:
 - After PSSP Version 3.4 became available, pSeries servers were introduced. They are high-end 4-, 8-, 16-, or 32-way GB processor POWER4 servers. The physical components can be assigned to separate logical partitions (LPARs) within one physical server. You can have up to sixteen LPARs. Each LPAR functions as an individual node within that frame and each is fully functional as a PSSP node in a Cluster 1600 or SP system. With PSSP Version 3.5, these nodes require AIX 5L Version 5.1. pSeries servers are supported in an SP-attached configuration with the SP Switch2, the SP Switch, or no switch, or in a clustered server configuration.
 - The IBM 375/450 MHz POWER3 SMP Thin Node and Wide nodes have been available with 375 MHz processors in an SP frame. Now you can have these thin and wide nodes with 450 MHz processors.
- ▶ Switch
 - With the control workstation running PSSP Version 3.4 (or later) software, you have optional switch connectivity. This means you can use the SP Switch2 with newer nodes and still keep older nodes in your system. Any node not supported on the SP Switch2 can remain in the system, but not connected to the SP Switch2.
 - For both the SP Switch and SP Switch2, you can specify which nodes you want to exclude from serving as a switch primary or primary backup, and which nodes are available for that purpose.
- ▶ SP-attached servers

SP system partitioning is supported by default in switchless systems with at least one SP node frame. For each SP node frame in the system, there are 16 available switch port numbers. Each SP node and each SP-attached

server logical node in the system must be assigned a switch port number. The number of SP-attached servers that can be added to the system is limited by the number of available switch port numbers.

- ▶ PSSP use of 64-bit kernel
 - PSSP Version 3.5 and the software stack it supports, including General Parallel File System (GPFS), Parallel Environment Message Passing Interface (MPI), Parallel Engineering Scientific Subroutine Language (Parallel ESSL), Low-Level Application Programming Interface (LAPI), Kernel Low-Level Application Programming Interface (KLAPI), and IBM Virtual Shared Disks, now provide support for use of 64-bit kernels. This support is essential for running applications with large real memory requirements. 64-bit kernel support only exists on those control workstations and nodes that meet the following minimal requirements:
 - PSSP Version 3.5 or later
 - AIX 5L Version 5.1 or later service level
 - The 64-bit kernel on hardware that AIX supports use of the 64-bit kernel

14.2.3 Hardware requirements

PSSP must be run on a Cluster 1600, or the equivalent cluster building blocks without the Cluster 1600 designation. These include:

- ▶ A control workstation. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*, GA22-7281 for the recommended hardware configuration.
- ▶ An RS/6000 SP system.
- ▶ One to 64 pSeries RS/6000 servers, depending on type.
- ▶ If you have non-SP cluster servers, there are cabling requirements that you must also consider. Refer to *IBM @server Cluster 1600 Hardware Planning, Installation, and Service*, GA22-7863 for more information.

14.2.4 Software requirements

The minimum PSSP software requirements for your RS/6000 SP system or Cluster 1600 system are as follows:

- ▶ AIX 5L Version 5.1 Base Operating System. AIX service is required at the 5100-03 Recommended Maintenance package (APAR IY32749 for service orders).
- ▶ PSSP Version 3.5.

- ▶ At least one of the following:
 - C for AIX Version 6.0 or later.
 - VisualAge® C++ Professional Version 6.0 for AIX or later.
 - PSSP must run on all partitions in a pSeries cluster.
 - PSSP Version 3.5 contains IBM VisualAge C++ Professional for AIX, Version 50 Runtime Modules.

14.3 General Parallel File System for AIX 5L

General Parallel File System (GPFS) for AIX 5L Version 2.1, a shared file system for serial and parallel applications, provides a single global file system for multiple nodes or servers, and is ideal for high-performance parallel file transfer and parallel I/O to single or multiple files.

This new version of GPFS for AIX 5L has the following enhancements:

- ▶ The 64-bit kernel is supported.
- ▶ High Availability Cluster Multi-Processing (HACMP) is no longer required to run GPFS in a non-SP or non-Parallel System Support Programs (non-PSSP) environment.
- ▶ It is supported with AIX 5L only.

14.3.1 Product positioning

GPFS is particularly appropriate in an environment where the aggregate peak need for data exceeds the capability of a distributed file system server.

- ▶ GPFS is beneficial to cluster customers who need a global UNIX file system to move data quickly and with high availability among nodes, and into and out of the node cluster. These customers are found in the S&TC, Business Intelligence, and Server Consolidation solution areas.
- ▶ GPFS enables UNIX customers to participate in the growth of clustering without the performance impediment implicit in single-system file systems such as NFS, Distributed File System (DFS™), and Journaled File System (JFS).
- ▶ Compute and data clustering, as pioneered by the RS/6000 SP, and extended to workstation and server clusters, offers value to UNIX customers by consolidating resources for easier system management. However, file systems such as NFS, DFS, and JFS, do not perform well in multiple system environments, and performance worsens as the number of systems trying to share the file system increases. This inhibits customer growth in the use of

clusters. GPFS overcomes this inhibitor with performance designed for the cluster environment.

The strengths of GPFS

GPFS is a powerful file system offering:

- ▶ Improved system performance
- ▶ Assured file consistency
- ▶ High recoverability and increased data availability
- ▶ Enhanced system flexibility
- ▶ Simplified administration

Improved system performance

Using GPFS to store and retrieve your files can improve system performance by:

- ▶ Allowing multiple processes or applications on all nodes in the nodeset simultaneous access to the same file using standard file system calls.
- ▶ Increasing aggregate bandwidth of your file system by spreading reads and writes across multiple disks.
- ▶ Balancing the load evenly across all disks to maximize their combined throughput. One disk is no more active than another.
- ▶ Supporting large amounts of data.
- ▶ Allowing concurrent reads and writes from multiple nodes. This is a key concept in parallel processing.

Assured file consistency

GPFS uses a sophisticated token management system to provide data consistency while allowing multiple independent paths to the same file by the same name from anywhere in the cluster. Even when nodes are down or hardware resource demands are high, GPFS can find an available path to file system data.

High recoverability and increased data availability

GPFS is a logging file system that creates separate logs for each node. These logs record the allocation and modification of metadata aiding in fast recovery and the restoration of data consistency in the event of node failure. GPFS failover support allows you to organize your hardware into a number of failure groups to minimize single points of failure. A failure group is a set of disks that share a common point of failure that could cause them all to become simultaneously unavailable. In order to assure file availability, GPFS maintains each instance of replicated data on disks in different failure groups.

Enhanced system flexibility

With GPFS, your system resources are not frozen. You can add or delete disks while the file system is mounted. When the time is right and system demand is low, you can rebalance the file system across all currently configured disks. You can also add new nodes without having to stop and restart the GPFS daemon.

Simplified administration

GPFS commands save configuration and file system information in one or more files, collectively known as GPFS cluster data. The GPFS administration commands are designed to keep these files synchronized between each other and with the GPFS system files on each node in the nodeset, thereby ensuring accurate configuration data. Most GPFS administration tasks can be performed from any node running GPFS.

14.3.2 Hardware requirements

In the cluster environment, the following are the basic hardware requirements:

- ▶ An existing system cluster from which to define a GPFS cluster
- ▶ Up to 1,024 external shared disks or disk arrays (including Serial Storage Architecture, Enterprise Storage Server, and IBM TotalStorage FASt500) with the correct adapters
- ▶ An IP network of sufficient bandwidth (minimum of 100 Mb per second) for GPFS daemon communications (token management, and so forth)

In the SP environment, the following are the basic hardware requirements:

- ▶ An RS/6000 SP
- ▶ Enough disk space to contain the file system (up to 1,024 logical devices)

Additional hardware is required to operate in a switch environment:

- ▶ SP Switch or SP Switch2
- ▶ SP Switch MX, MX2, or PCI-X adapter to exploit Low-Level Application Programming Interface (LAPI) as the communication protocol

Additional hardware is required to operate in a direct-attached disk environment:

- ▶ Up to 1,024 external shared disks or disk arrays (including Serial Storage Architecture, IBM Enterprise Storage Server, and TotalStorage FASt500) with the correct adapters
- ▶ An IP network of sufficient bandwidth (minimum of 100 Mb per second) for GPFS daemon communications

14.3.3 Software requirements

The following are the software requirements:

- ▶ AIX 5L Version 5.1 or later modification level (5765-E61) and APAR IY33002

On an RS/6000 SP or Cluster 1600 with a Switch:

- ▶ Parallel System Support Programs (PSSP) for AIX Version 3.5 or later (5765-D51)
- ▶ The IBM Virtual Shared Disk image and Recoverable Virtual Shared Disk image of PSSP

14.4 AIX Workload Manager

AIX Workload Manager (WLM) is an operating system feature introduced in AIX Version 4.3.3 and enhanced in AIX 5L. It is part of the operating system kernel at no additional charge. Provides a policy-based method for managing system workload and system resources.

Current UNIX offerings for partitioning and workload management have clear architectural differences. Partitioning creates isolation between multiple applications running on a single server, hosting multiple instances of the operating system. Workload management supplies effective management of multiple, diverse workloads to efficiently share a single copy of the operating system and a common pool of resources. AIX Workload Manager is a workload management feature.

Customers, system administrators, performance consultants, and managers should be aware that Workload Manager is not a tuning tool. AIX WLM is a resource management tool that specifies the relative importance of each workload by classes, tiers, limits, shares, and rules.

WLM is ideally suited to balance the demands or requests of competing workloads when one or more resources are constrained. It prevents a relatively uncontrolled way of resource scheduling for different applications on the system.

AIX Workload Manager includes the following capabilities:

- ▶ WLM allows the system administrator to divide resources between jobs without having to partition the system.
- ▶ Provides isolation between user communities with very different system behaviors. This can prevent effective starvation of workloads with certain characteristics, such as interactive or low CPU usage jobs, by workloads with other characteristics, such as batch or high memory usage jobs.

- ▶ The setup of WLM is much simpler than partitioning, where reinstallation and reconfiguration are required. With WLM, a single operating system manages the entire system and all jobs, so only one system has to be administered.
- ▶ WLM manages percentages of CPU time rather than CPUs. This allows control over CPU resources at a finer granularity.
- ▶ CPU time, memory, and I/O bandwidth are managed separately. Therefore, different styles of applications can be managed.
- ▶ AIX WLM delivers the basic ability to give system administrators more control over how scheduler, Virtual Memory Manager (VMM), and device driver calls allocate CPU, physical memory, and I/O bandwidth to classes-based users, groups, application paths, process types, or application tags. It allows a hierarchy of classes to be specified, processes to be automatically assigned to classes by their characteristics, and manual assignment of processes to classes.
- ▶ Administrators are spared from writing complex scripts.
- ▶ 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.

- ▶ These capabilities can be easily managed through Web-based System Manager, SMIT, shell scripts, or command line interfaces. WebSM enables management of AIX systems on the Internet from anywhere using an intuitive, object-oriented, easy-to-use GUI. WebSM will automatically show the current resource utilization by class whenever the top level of WLM management is shown in WebSM.

The following outlines the enhancements that the latest version of AIX WLM offers over earlier releases:

- ▶ Disk I/O Bandwidth, a new resource type, is introduced in addition to existing resources such as CPU cycles and real memory.
- ▶ An Application Programming Interface (API) enables external applications to modify system behaviors.
- ▶ System administrators can manually reclassify processes independent of the classification rules.
- ▶ This function enables multiple instances of the same application to exist in different classes. Using Application Tag API, applications can enable automatic assignment of multiple instances of the same application in different classes.
- ▶ New subclasses add ten times the granularity of control (from 27 classes to 270 controllable classes).
- ▶ System administrators can delegate subclass management to users or groups.
- ▶ Now fully dynamic, Workload Manager allows an entire configuration to be changed while it is running.
- ▶ Application path name extends wild card flexibility to user name and group name.
- ▶ The accounting subsystem, a new feature, allows users to perform resource usage accounting per WLM class in addition to the standard accounting per user or group.
- ▶ The AIX accounting system utility allows system administrators to collect and report the use of various system resources by user, group, or WLM class.
- ▶ When process accounting is turned on, AIX records statistics about the process resource usage in an accounting file when the process exits. This accounting record now includes a 64-bit numeric key representing the name of the WLM class, and the process belonged to.
- ▶ The accounting command, **acctcom**, allows the display of process resource usage statistics per user, group, or WLM class.

14.5 AIX Performance ToolBox and Performance AIDE

The Performance Toolbox (PTX®) is a comprehensive tool set for monitoring and analyzing system performance in local and distributed environments. In conjunction with the Performance AIDE, PTX enables you to view both live and recorded data.

The AIDE also provides agents for creating 24 x 7 recordings of large sets of performance metrics, and for filtering data based on user-customized criteria. Filtering criteria can be used to generate events to a monitoring console or execute administration scripts. Additional utilities are provided to convert recorded data into formats suitable for import by third-party spreadsheets. By providing an umbrella for tools that can be used to analyze performance data and control system resources, PTX assists the system administrator in keeping track of available tools and in applying them in appropriate ways. This is done through a customizable menu interface. Tools can be added to menus, either with fixed parameters to match specific situations or in a dialog window.

AIX Performance Toolbox and AIX Performance AIDE Version 3.1 delivers the following features and benefits:

- ▶ Full support for AIX 5L Versions 5.1 and 5.2.
- ▶ New Jtopas analysis interface for viewing top resource consumers. The Jtopas application creates simplified snapshot views of overall activity for near real-time and playback analysis.
- ▶ New reporting capabilities that support hourly and daily usage.
- ▶ Enhanced agent support for top resource analysis.
- ▶ Support for aggregated adapter metrics.
- ▶ Support for new system metrics.

AIX Performance AIDE Version 3.1 enhancements:

- ▶ Common agent for AIX 5L Version 5.1 and AIX 5L Version 5.2.
- ▶ Agent supports new optimized top resource collection and recording.
- ▶ Top agent provides continuous recording function for longer term analysis and playback.
- ▶ Collection framework provides new performance metrics available in AIX 5L and DLPAR environments.
- ▶ Performance AIDE maintains an unlimited node license.

14.5.1 Hardware requirements

The following are the basic hardware requirements:

- ▶ RS/6000 or pSeries system with a minimum of 256 MB memory
- ▶ Graphics adapter for Performance Toolbox user interface operation
- ▶ Network adapter needed for Performance Toolbox remote monitoring

14.5.2 Software requirements

AIX 5L Version 5.1 with 5100-01 Recommended Maintenance package (included on the September 2001 or later AIX Update CD) or later is the required software level.

14.6 AIX Fast Connect for POWER Version 3.1

AIX Fast Connect delivers outstanding performance and attractive price for any number of Windows clients.

Using Fast Connect, PCs can take advantage of the performance, scalability, and reliability of AIX. Multiple PC servers can be consolidated onto a single, scalable IBM *@server* pSeries or RS/6000, to reduce system complexity and simplify systems management.

Fast Connect is internationalized, which includes messages in 26 languages supported by AIX and Unicode capability to support international characters in file and directory names and other character strings.

With Fast Connect, AIX file and print services are available to Windows 95/98/NT/2000 users as part of their Network Neighborhood with no additional software required. Windows clients can access AIX files using the Journaled File System (JFS), JSF2, CD File System (CDFS), Network File System (NFS) mounted subsystems, and DCE/DFS file systems, and AIX printing services using Microsoft's Common Internet File System (CIFS) and Server Message Block (SMB) protocols over TCP/IP.

IBM has one of the strongest UNIX-PC interoperability records in the industry. IBM has demonstrated leadership in understanding and managing the complexity of heterogeneous environments by:

- ▶ Optimizing the use of shared resources
- ▶ Leveraging the advantages of UNIX for PC LANs connecting clients with enterprise servers
- ▶ Supporting a broad range of IBM and third-party products



Cluster systems

Today's globally networked economy demands powerful information systems that are available to customers, business partners, and employees around the clock and around the world. In today's fiercely competitive environment, high availability is important because it minimizes the risk of an outage, and increases the availability of the system. A system outage can be costly in terms of lost productivity, lost revenue, and lost customers. Moreover, it also takes years to gain a good solid reputation and only seconds to tarnish it. With the increasing acceptance of e-business, companies cannot afford to risk impacting customer service levels due to systems outages, both planned and unplanned. The demands of the customers and end users require round the clock system access. These are the factors driving the high availability phenomena.

High availability views availability not as a series of replicated physical components, but rather as a set of system-wide, shared resources that cooperate to guarantee 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.

In this chapter, we will discuss high availability solutions as follows:

- ▶ HACMP (High Availability Cluster Multi-Processing) for AIX
- ▶ HAGEO (High Availability Geographic Cluster) for AIX
- ▶ Cluster 1600

15.1 HACMP

This section describes the HACMP solution from a functional perspective.

15.1.1 Basic concepts of HACMP

Clustering is the linking of two or more computers or nodes into a single, unified resource. High availability clusters are designed to provide continuous access to business-critical data and applications through component redundancy and application failover.

High Availability Cluster Multi-Processing (HACMP) is IBM's software for building highly available clusters on IBM Scalable POWERParallel systems or a combination of pSeries systems. It is supported by a wide range of IBM @server pSeries systems, including the p690, storage systems, and network types. HACMP builds on the inherent reliability of the hardware to provide greater uptime for applications and enables upgrades and reconfiguration without interrupting operations. These servers or LPARs can also be part of an IBM @server Cluster 1600, which can simplify multisystem management and help to reduce cost of ownership.

HACMP provides concurrent access to IT resources and the fault resilience required for business-critical applications. It is designed to automatically detect system or network failures and eliminates a single point-of-failure by managing failover to a recovery processor with a minimal loss of end-user time.

HACMP can also detect software problems that are not severe enough to interrupt proper operation of the system, such as process failure or exhaustion of system resources. HACMP monitors, detects, and reacts to such failure events, allowing the system to stay available during random, unexpected software problems. HACMP can be configured to react to hundreds of system events.

HACMP makes use of redundant hardware configured in the cluster to keep an application running, restarting it on a backup processor if necessary. This minimizes expensive downtime for both planned and unplanned outages and provides flexibility to accommodate changing business needs. Up to 32 pSeries or IBM RS/6000 servers can participate in an HACMP cluster, which is ideal for an environment requiring horizontal growth with rock-solid reliability.

The difference between fault tolerance and high availability is a fault tolerant environment has no service interruption, while an environment with high availability has a minimal service interruption. Many customers are willing to accept a small amount of down time with high availability rather than pay the much higher cost of providing fault tolerance. Additionally, in most highly

available configurations, the backup processors are available for use during normal operation.

15.1.2 Components

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 pSeries uniprocessor, 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.

- ▶ Service network adapter

The service network adapter is the primary connection between an HACMP node and a network. A node can have one or more service network adapter labels for each physical network to which it connects. This adapter label is used for cluster TCP/IP traffic. Its address is published by the Cluster Information Program (Clinfo) to application programs that want to use cluster services.

- ▶ Standby network adapter

The standby network adapter backs up a service network adapter. The service network adapter can be on the local node or, if IP address takeover is enabled, on a remote node. If a service network adapter on the local node fails, the HACMP software swaps the standby network address with the service network address. If the local node is designated to take over the network address of a peer node should that node fail, the standby network adapter label on the local node assumes the IP address of the service network adapter on the failed node.

- ▶ Boot adapters and IP address takeover

IP address takeover (IPAT) is an HACMP facility that lets one node acquire the network address of another cluster node. For IPAT to work correctly, however, you must configure the boot adapter to the AIX TCP/IP configuration and then configure both the boot and service IP addresses as part of the HACMP adapter configuration. Both addresses must be assigned to the same network. Cluster nodes use the boot label after a system reboot and before the HACMP software is started, or after it is stopped gracefully with or without takeover. When a node is forced down, however, the adapter does not revert to its boot address. When HACMP is started on a node, the node's service adapter is reconfigured to use the service label (address) instead of the boot label. If the node should fail, a takeover node acquires the failed node's service address on its standby adapter, making the failure transparent to clients using that specific service address.

During the reintegration of the failed node, which comes up on its boot address, the takeover node will release the service address it acquired from the failed node. Afterwards, the reintegrating node will reconfigure its boot address to its reacquired service address. It is important to know that the boot address does not use a separate physical adapter, but instead is a second name and IP address associated with a service adapter.

► Persistent node address

The persistent node address (persistent IP label) is an IP alias that can be assigned to an HACMP network on a specified node. Assigning a persistent IP label to a network on a node allows you to have a node-bound address on a cluster network that you can use for administrative purposes to access a specific node in the cluster, using the **ping** or **telnet** commands.

A persistent IP label is an address that:

- Always stays on the same node (is node-bound)
- Co-exists on an adapter that already has a service or boot label defined
- Does not require installing an additional physical adapter on that node
- Is not part of any resource group

Once a persistent IP label is configured on a particular network on a particular node, it becomes available on that node at boot time and remains configured on that network when HACMP is shut down on that node.

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 15-1 shows an example of HACMP configuration.

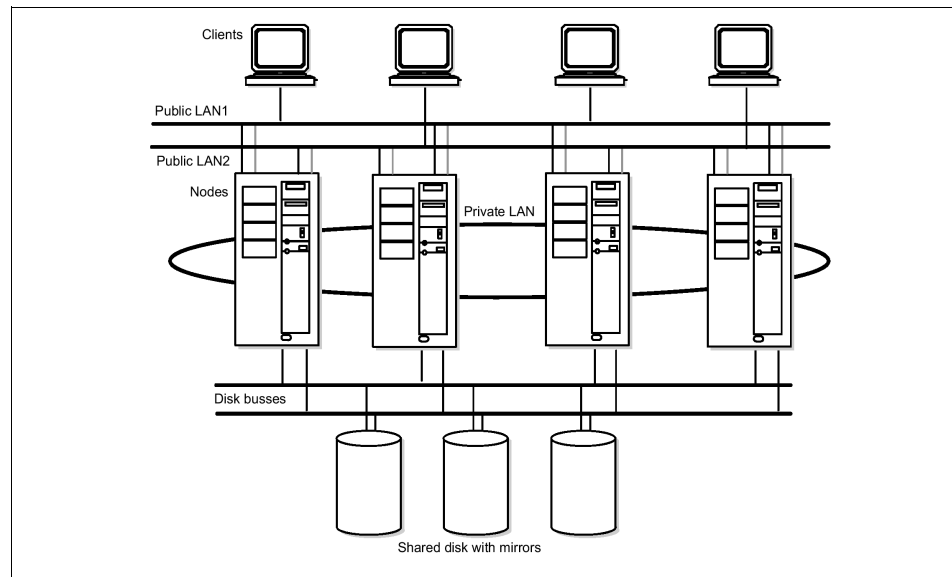


Figure 15-1 An example of high availability cluster

15.1.3 Resource groups

HACMP provides a highly available environment by:

- ▶ Identifying a set of cluster-wide resources essential to uninterrupted processing
- ▶ Defining relationships among nodes that ensure these resources are available to client processes.

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 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.

HACMP considers the following as resource types:

- ▶ Volume groups
- ▶ Disks
- ▶ File systems
- ▶ File systems to be NFS mounted
- ▶ File systems to be NFS exported
- ▶ Service IP addresses
- ▶ Applications

Each resource in a cluster is defined as part of a resource group. This allows you to combine related resources that need to be together to provide a particular service. A resource group also includes the list of nodes that can acquire those resources and serve them to clients.

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. You define the takeover relationship of a resource group by assigning it one of the following type designations:

- ▶ Cascading (with or without the Cascading without Fallback attribute)
- ▶ Rotating
- ▶ Concurrent

Each of these types describes a different set of relationships between nodes in the cluster, and a different set of behaviors upon nodes entering and leaving the cluster.

Cascading resource groups

All nodes in a cascading resource group are assigned priorities for that resource group. These nodes are said to be part of that group's resource chain. In a cascading resource group, the set of resources cascades up or down to the highest priority node active in the cluster. When a node that is serving the resources fails, the surviving node with the highest priority takes over the resources.

A new parameter called Cascading Without Fallback (CWOFF) which is a new feature in HACMP Version 4.4.1, is an attribute of cascading resource groups that defines its fallback behavior. With this option enabled, a cascading resource group will not fallback to a higher priority node as it joins or reintegrates into the cluster. A cascading group with CWOFF set to FALSE will exhibit fallback behavior.

You can use cascading configuration in case you have servers (nodes) in a different configuration, for example, the first node has better performance and many more resources, such as memory, processors and so on, than the second node. In this case, you can use the slowest, less powerful node as a temporary server for your resources and applications if the first node fails. If the first node comes up, resources automatically will be moved to the first node with highest priority.

Rotating resource groups

A rotating resource group is associated with a group of nodes, rather than a particular node. A node can be in possession of a maximum of one rotating resource group per network. As participating nodes join the cluster for the first time, they acquire the first available rotating resource group per network until all the groups are acquired. The remaining nodes maintain a standby role.

When a node holding a rotating resource group leaves the cluster, either because of a failure or gracefully while specifying the takeover option, the node with the highest priority and available connectivity takes over. Upon reintegration, a node remains as a standby and does not take back any of the resources that it had initially served.

Concurrent resource groups

A concurrent resource group may be shared simultaneously by multiple nodes. The resources that can be part of a concurrent resource group are limited to volume groups with raw logical volumes, raw disks, and application servers.

When a node fails, there is no takeover involved for concurrent resources. Upon reintegration, a node again accesses the resources simultaneously with the other nodes. As an example of using concurrent resource group is a database, where you can spread their workload across the cluster.

15.1.4 Cluster configurations

This section provides the types of cluster configurations supported by the HACMP software.

Non-concurrent disk access configurations

The possible non-concurrent disk access configurations are:

- ▶ Hot-standby
- ▶ Rotating standby
- ▶ Mutual takeover
- ▶ Third-party takeover

Hot-standby configuration

Figure 15-2 illustrates a two node cluster in a hot-standby configuration.




Figure 15-2 Hot-standby configuration

In this configuration, there is one cascading resource group consisting of disks and their constituent volume groups and file systems. Node 1 has a priority of 1 for this resource group, while node 2 has a priority of 2. During normal operations, node 1 provides all critical services to end users. Node 2 may be idle or may be providing non-critical services, and hence is referred to as a hot-standby node. When node 1 fails or has to leave the cluster for a scheduled outage, node 2 acquires the resource group and starts providing the critical services.

The advantage of this type of a configuration is that you can shift from a single-system environment to an HACMP cluster at a low cost by adding a less powerful processor. Of course, this assumes that you are willing to accept a lower level of performance in a failover situation. This is a trade-off that you will have to make between availability, performance, and cost.

Rotating standby configuration

This configuration is the same as the previous configuration except that the resource groups used are rotating resource groups. In the hot-standby configuration, when node 1 reintegrates into the cluster, it takes back the resource group since it has the highest priority for it. This implies a break in service to the end users during reintegration. If the cluster is using rotating resource groups, reintegrating nodes do not reacquire any of the resource groups. A failed node that recovers and rejoins the cluster becomes a standby node. You must choose a rotating standby configuration if you do not want a break in service during reintegration. Since takeover nodes continue providing services until they have to leave the cluster, you should configure your cluster with nodes of equal power. While more expensive in terms of CPU hardware, a rotating standby configuration gives you better availability and performance than a hot-standby configuration.

Mutual takeover configuration

Figure 15-3 on page 636 illustrates a two node cluster in a mutual takeover configuration.




Figure 15-3 Mutual takeover configuration

In this configuration, there are two cascading resource groups: A and B with one volume group each. Node 1 has priorities of 1 and 2 for resource groups A and B respectively, while Node 2 has priorities of 1 and 2 for resource groups B and A respectively. During normal operations, nodes 1 and 2 have control of resource groups A and B respectively, and both provide critical services to end users. If either node 1 or node 2 fails, or has to leave the cluster for a scheduled outage, the surviving node acquires the failed node's resource groups and continues to provide the failed node's critical services.

When a failed node reintegrates into the cluster, it takes back the resource group for which it has the highest priority. Therefore, even in this configuration, there is a break in service during reintegration. Of course, if you look at it from the point of view of performance, this is the best thing to do, since you have one node doing the work of two when any one of the nodes is down.

Third-party takeover configuration

Figure 15-4 on page 637 illustrates a three node cluster in a third-party takeover configuration.




Figure 15-4 Third-party takeover configuration

This configuration can avoid the performance degradation that results from a failover in the mutual takeover configuration. Here the resource groups are the same as the ones in the mutual takeover configuration. Also, similar to the previous configuration, nodes 1 and 2 each have priorities of 1 for one of the resource groups, A or B. The only thing different in this configuration is that there is a third node that has a priority of 2 for both the resource groups.

During normal operations, node 3 is either idle or is providing non-critical services. In the case of either node 1 or node 2 failing, node 3 takes over the failed node's resource groups and starts providing its services. When a failed node rejoins the cluster, it reacquires the resource group for which it has the highest priority.

So, in this configuration, you are protected against the failure of two nodes and there is no performance degradation after the failure of one node.

Concurrent disk access configurations

A concurrent disk access configuration usually has all its disk storage defined as part of one concurrent resource group. The nodes associated with a concurrent resource group have no priorities assigned to them.

In the case of a node failure, a concurrent resource group is not explicitly taken over by any other node, since it is already active on the other nodes. However, in order to somewhat mask a node failure from the end users, you should also have cascading resource groups, each containing the service IP address for each node in the cluster. When a node fails, its service IP address will be taken over by another node and users can continue to access critical services at the same IP address that they were using before the node failed.

The concurrent access feature enhances the benefits provided by an HACMP cluster. Concurrent access allows from two to eight processors to simultaneously access a database or applications residing on shared external disks. Using concurrent access, a cluster can offer nearly continuous availability of resources that rivals fault tolerance, but at much lower cost. Additionally, concurrent access provides higher performance, eases application development, and allows horizontal growth.

15.1.5 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, Version 2.1, Version 3, Version 3.1, and Version 3.1.1, all of which were based on AIX Version 3.2.5. Since this time, there have been five releases of HACMP Version 4, which are 4.1, 4.2, 4.3, 4.4, and 4.5. 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. We discuss a brief history of HACMP versions in 15.1.6, “HACMP history” on page 640.

Scalability, support of large clusters, and, therefore, large configurations of nodes including 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, that allow disks to be accessed by a large number of systems without a performance penalty.

If it is crucial to have constant shared access with no time lost to failovers or fallbacks, you may want to choose a concurrent configuration. To define a concurrent configuration, you must have the HACMP/ES Concurrent Resource Manager (ESCRM) software installed and the application must support concurrent access.

The following sections will describe a brief description of HACMP and HACMP/ES and also the features that comprise the HACMP and HACMP/ES. The following options are referred to as:

- ▶ HACMP classic
 - High Availability Subsystem (HAS)
 - Concurrent Resource Manager (CRM)
 - High Availability Network File System (HANFS); this is included in HACMP and HACMP/ES since Version 4.4.0
- ▶ HACMP/ES
 - Enhanced Scalability (ES)
 - Enhanced Scalability Concurrent Resource Manager (ESCRM)

Overall features and benefits of HACMP are summarized in Table 15-1.

Table 15-1 Features and benefits of HACMP

Feature	Benefits
Horizontal scalability	<ol style="list-style-type: none"> 1. Boosts overall performance and capacity by sharing disk resources of clustered systems, thereby spreading applications across pSeries and RS/6000 servers 2. Provides scalable growth with reduced investment and increased system availability 3. Enables mix of multiprocessor nodes for application performance and disk sharing
Support for multiple availability configurations	<ol style="list-style-type: none"> 1. Allows tailoring for a flexible solution that can change with your business 2. Allows growth for availability in either standby or mutual takeover modes in 2- to 32-node clusters
Breadth of platform support	<ol style="list-style-type: none"> 1. Runs across the IBM @server pSeries and RS/6000 product lines 2. Supports the IBM @server Cluster 1600 offerings

Feature	Benefits
Reliable scalable cluster technology	<ol style="list-style-type: none"> 1. Detects and reacts to hundreds of system events to protect the software infrastructure 2. Provides unrivaled levels of availability using standard components 3. Allows determination of the best node to take over after a failure 4. Further enhances availability through application monitoring and user-defined events 5. Allows applications to better integrate with HACMP
Support for cluster management SMIT facilities	<ol style="list-style-type: none"> 1. Provides common interfaces to install and configure highly available cluster systems as well as maintain them on the network 2. Provides comprehensive, highly available monitoring utilities to manage and tune clusters 3. Enables visualization of relationships between cluster hardware and software resources 4. Allows management of clusters within an enterprise environment
Compatibility with previous HACMP versions	<ol style="list-style-type: none"> 1. Reuses configurations to take advantage of current customer failover control scenarios to provide node-at-a-time upgrades 2. Makes transition easier, removing the necessity to reinstall or reconfigure clusters when upgrading

15.1.6 HACMP history

This section provides a historical view of HACMP Version 1.1 through 4.3.1. Table 15-2 shows HACMP life cycle information.

Table 15-2 HACMP life cycle information

Program number	VRM	Announced	Marketing withdrawn	Service discontinued
5765-A86	4.2.2	10/31/1997	03/01/2000	01/31/2001
5765-D28	4.3.0	10/05/1998	07/23/1999	12/31/2001
5765-D28	4.3.1	07/20/1999	01/01/2001	12/31/2001
5765-E54	4.4.0	06/20/2000	07/13/2001	12/21/2003
5765-E54	4.4.1	06/26/2001	06/30/2003	12/31/2003

Program number	VRM	Announced	Marketing withdrawn	Service discontinued
5765-E54	4.5.0	06/25/2002	-	12/31/2004
5765-F62	5.1.0	06/24/2003	-	-

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.

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, cllockd, 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 is loosely coupled multiprocessing, which allowed two processors to work on the same shared workload.

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).

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.

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.

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 offline.

- ▶ The CRM feature optionally adds concurrent shared access management for supported disks.
- ▶ Support is added for IBM's symmetric multiprocessor (SMP) machines.

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).

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.

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.

HACMP ES Version 4.2.1

The first release of HACMP ES, Version 4.2.1, was announced shortly after HACMP Version 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 Version 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.

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 `clidare` 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 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.

HACMP and HACMP ES Version 4.3

Custom snapshot methods were 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 (Run `smi t hacmp` and select **Cluster System Management** → **Task Guide For Creating a Shared Volume Group**).
- ▶ Multiple pre- and post-event scripts are 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 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.

HACMP and HACMP ES Version 4.3.1

Support for target mode SSA was 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 (run `smi t hacmp` and select **Cluster System Management** → **Swap Adapter**).
- ▶ The location of HACMP log files can now be defined by the user by specifying an alternative path name using SMIT (run `smi t hacmp` and select **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.

15.1.7 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 is designed to detect system failure and manage failover to a recovery processor with a minimal loss of end-user time. 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.

15.1.8 HACMP and HACMP ES Version 4.4.1

This section describes general information on HACMP Version 4.4.1.

Enhancements of HACMP Version 4.4.1

HACMP and HACMP ES Version 4.4.1 was introduced in July 2001 with extended capabilities as follows:

- ▶ Enhanced failover support, including the ability to handle combinations of multiple failures in certain circumstances
- ▶ Improved usability, including enhanced SMIT Interface for Configuring Volume Groups and Networks, customized pager notifications, and HAGEO site configuration data
- ▶ Use of C-SPOC to replace a failed disk
- ▶ Expanded device support, including Multiple Logical Interfaces on the same ATM Network Adapter, OEM Disk API, and 32 Nodes in Concurrent Access Mode Volume Groups
- ▶ Replacement of hot-plug capable PCI Network Adapters
- ▶ New Highly Available Network Service demonstration code (sample scripts)

Moreover, HACMP ES Version 4.4.1 introduces more enhancements such as:

- ▶ Enhanced SMIT interface for configuring volume groups: Automatic discovery and import of shared volume groups
- ▶ Enhanced SMIT interface for configuring networks: Automatic discovery of IP addresses
- ▶ Support of multiple logical interfaces on the same ATM network adapter: ATM Classic IP and ATM LANE
- ▶ Replacement of a hot-plug PCI Network Adapter through SMIT interface
- ▶ SMIT interface for custom pager notification in response to a cluster event
- ▶ Improved selection of performance tuning parameters
- ▶ Support for up to 32 nodes in concurrent access mode volume groups
- ▶ Ability to mount all file systems for a resource group
- ▶ Extended Inactive Takeover support for cascading resource groups with more than two nodes

- ▶ User-defined events improvements
- ▶ Dynamic node priority policies: Takeover node defined by resource variable
- ▶ Dynamic, enhanced handling of resource groups
- ▶ Other HACMP enhancements
- ▶ Enhanced failover support, including handling new combinations of failures
- ▶ Enhanced clverify functionality
- ▶ Summary of Events in hacmp.out with HTML readability
- ▶ OEM disk API enablement
- ▶ HAGEO integration: HACMP SMIT interface for site configuration
- ▶ Support for stopping cluster services with the Forced option
- ▶ C-SPOC user password support
- ▶ Highly Available Network Service demonstration code (sample scripts)
- ▶ RSCT support for HACMP/ES Star Topology Configuration

15.1.9 HACMP and HACMP ES Version 4.5

IBM customers with mission-critical environments can better protect their multiprocessing systems with the High Availability Cluster Multi-Processing Version 4.5 for AIX 5L introduced on July 12, 2002. HACMP 4.5 offers improved usability and performance, easier configuration, and additional hardware and software support for IBM @server Cluster 1600, IBM @server pSeries, and RS/6000 customers with mission-critical applications.

Enhancements of HACMP Version 4.5

New features of HACMP Version 4.5 include:

- ▶ Reduced Fallover Time through parallel processing of resources
- ▶ Integration with AIX Workload Manager (WLM) for better operational performance after fallover
- ▶ Streamlined configuration process using automated Network Discovery
- ▶ Improved security for cluster administration using HAView and HATivoli
- ▶ Persistent IP address support for applications (such as Tivoli), which require invariant node addresses
- ▶ Expanded WAN and X.25 support
- ▶ Cluster status information in a Web Browser
- ▶ Enhanced Custom Pager Notification support

For HACMP Enhanced Scalability (ES) and ES/Concurrent Resource Manager (ESCRM), the following additional enhancements are provided:

- ▶ Easier configuration of AIX Enhanced Concurrent Mode to define concurrent volume groups on any disk device
- ▶ New Application Availability Analysis Tool for customer measurement of system and application up time
- ▶ Tighter integration with IBM Generic Parallel File System (GPFS) Version 1.5 Cluster Filesystem
- ▶ Monitoring and recovery from loss of volume group quorum
- ▶ IP Aliasing support for multiple applications on each network adapter
- ▶ 64-bit Client Info API for user-written cluster-aware applications
- ▶ User-specified Resource Groups Processing Order

15.1.10 HACMP Version 5.1

IBM customers can better protect their critical business applications from failures with the world-wide capabilities of IBM High Availability Cluster Multi-Processing for AIX 5L Version 5.1.0 (HACMP Version 5.1).

HACMP Version 5 includes all of the functions previously included in HACMP Version 4, HAS, ES, CRM, and ESCRM. This simplification enables you to easily take advantage of, and achieve all of the benefits of all of the traditional flavors and powers of HACMP.

HACMP Version 5 provides base services for cluster membership, system management, configuration integrity and control, failover, and recovery for up to 32 servers or nodes. Easy-to-use status and monitoring facilities are included. Scalability functionality provides these capabilities across entire clusters, and allows customers to define their own HACMP events and monitor their applications. HACMP Version 5 fully supports administration of AIX 5L Enhanced Concurrent Mode, thus providing you concurrent shared-access management for all supported disk subsystems.

Now, HACMP Version 5.1.0 offers not only improved usability, functionality, and performance, but also a new single optional package that includes both ESS PPRC and HAGEO IP-based support for geographically dispersed data backup and disaster recovery. Customers now have a single, world-class source of protection for their mission-critical applications.

Enhancements of HACMP Version 5.1

New features of HACMP Version 5.1 include:

- ▶ Consolidation of all previous forms of HACMP (HAS, CRM, ES, and ESCRM) into one single valuable HACMP offering
- ▶ Fast Disk Fallover exploits AIX 5L Version 5.2 Enhanced Concurrent Mode, takes less than 10 seconds for disks to fall over, and minimizes application downtime
- ▶ Simplified installation and configuration, enabled through new streamlined user interface, reducing decision points
- ▶ Heartbeat by shared disk offers additional protection against network-partitioned cluster data divergence (split-brain syndrome)
- ▶ Enhanced security mechanism, removing need for /.rhosts Integrated Cluster File System utilizing General Parallel Files System (GPFS) Version 1.5
- ▶ Performance re-write of cluster verification and Cluster Single Point-Of-Control (C-SPOC) for faster fallover and increased administrative productivity
- ▶ Custom Resource Group Management facility for greater control of startup and fallback application behavior
- ▶ Extensive cluster information now provided in `c1stat` command output
- ▶ Optional HACMP/XD (Extended Distance) feature for ESS/PPRC and IP-based mirroring configurations
 - HACMP Version 5's new XD (Extended Distance) option provides automated data backup and disaster recovery across geographically-dispersed geographies, enabling you to protect business-critical applications and provide 24x7 service to your customers and staff.
 - HACMP/XD will immediately support IBM Enterprise Storage Subsystem (ESS) Peer-to-Peer Remote Copy (PPRC). This means HACMP/XD clusters now support automatic fallover of disks that are PPRC pairs, and creates a powerful disaster recovery solution for customers on ESS. HACMP/XD automates the management of PPRC, minimizes recovery time after an outage, and monitors your clustered environment to ensure mirroring of critical data is maintained at all times.
 - HACMP/XD IP-based mirroring will provide the well-known unlimited distance data mirroring of the current IBM HAGEO product. HACMP/XD will deliver a fully integrated copy of HAGEO Version 2.4, allowing a cluster of IBM @server pSeries or RS/6000 computers to be placed in two widely separated geographic locations, each maintaining an exact

replica of the application and data. Data synchronization during production, fallover, recovery, and restoration is provided.

15.1.11 Hardware supports matrix

HACMP works with IBM @server Cluster 1600, pSeries, and RS/6000 servers, and SP systems in a no-single-point-of-failure server configuration. HACMP supports the Cluster 1600, pSeries, and RS/6000 models designed for server applications and that meet the minimum requirements for internal memory, internal disk, and I/O slots.

System support

Table 15-3 shows pSeries and RS/6000 models and their corresponding upgrades that are supported in each version of HACMP.

Notes:

X Supported
Blank Not supported

Table 15-3 HACMP supported systems

Systems	Models	HACMP version		
		4.4.1	4.5	5.1
PCI Desktop Systems	Models 140, 150, 170, 240, 260, and 270	X	X	X
PCI Deskside Systems	Models E20, E30, F30, F40, F50, F80, 6F0, and 6F1 (pSeries 620)	X	X	X
Entry Systems	Models 25S, 250, and 25T	X	X	X
Compact Server Systems	Models C10 and C20	X	X	X
Desktop Systems	Models 370, 380, 390, 397, and 39H	X	X	X
Deskside Systems	Deskside Systems: Models 570, 57F, 580, 58F, 58H, 590, 59H, 591, 595, 7028-6E1 (pSeries 610), 7025-6F1 (pSeries 620), and 7028-6E4 (pSeries 630)	X	X	X
Rack Systems	Rack Systems: Models 98B, 98E, 98F, 990, 99E, 99F, 99J, 99K, B80, R10, R20, R21, R24, R50, R5U, S70, S7A, H10, H50, H70, H80, M80, 7028-6C1 (pSeries 610), 7028-6C4 (pSeries 630), 6H0, 6H1, 6M1 (pSeries 660), 7039-651 (pSeries 655), and 7038-6M2 (pSeries 650)	X	X	X

Systems	Models	HACMP version		
		4.4.1	4.5	5.1
High-End Servers	Models 681 (pSeries 690) and 671 (pSeries 670)		X	X
Symmetric Multiprocessor Server Systems	Models G30, J30, R30, R3U, G40, J40, R40, R4U, J50, R4U, S70, S7A, S80, and S85 (pSeries 680)	X	X	X
SP Systems	Models 204, 205, 206, 207, 208, 209, 20A, 2A4, 2A5, 2A7, 2A8, 2A9, 2AA, 304, 305, 306, 307, 308, 309, 30A, 3A4, 3A5, 3A7, 3A8, 3A9, 3AA, 3B4, 3B5, 3B7, 3B8, 3B9, 3BA, 404, 405, 406, 407, 408, 409, 40A, 500, 50H, 550, and 55H, including the 604 High Nodes, 604E High Nodes, the Power2 Super Chip (P2SC) nodes, and the 375 MHz POWER3 SMP Nodes	X	X	X

For a concurrent access configuration, HACMP Versions 4.4.1 and 4.5 requires one of the following devices.

- ▶ IBM 7131 SSA Multi-Storage Tower Model 405 (supports up to eight nodes; no CD-ROMs or tapes can be installed)
- ▶ IBM 7133 SSA Disk Subsystem Models 020, 600, D40, and T40 (supports up to eight nodes)
- ▶ IBM 7135 RAIDiant Array Models 110 and 210 (supports up to four nodes; dual controllers recommended)
- ▶ IBM 7137 Disk Array Subsystem Models 413, 414, 415, 513, 514, or 515 (supports up to four nodes)
- ▶ IBM 2105 Versatile Storage Server™ (VSS) Models B09 and 100 (supports up to four nodes)
- ▶ IBM 2102-F10 Fibre Channel RAID Storage Serve
- ▶ IBM Enterprise Storage Server (ESS) Models E10, E20, F10, and F20 (supports up to eight nodes using SCSI and Fibre Channel interfaces using IBM FC/FICON Features 3021, 3022 and 3023) - Only for HACMP Version 4.5

HACMP Version 5.1 supports a concurrent access configuration with all supported external storage systems.

Note: Certain non-IBM RAID systems can operate in concurrent I/O access environments. IBM will not accept Authorized Program Analysis Reports (APARs) if the non-IBM RAID offerings do not work properly with HACMP Version 4.5.0.

Device support

We will take a look at devices supported for each version of HACMP.

Communication adapters

HACMP supports the communication adapters listed in Table 15-4 on page 655.

Notes:

X	Supported
Blank	Not supported

Table 15-4 Communication adapters

Type	Feature code	Description	HACMP version		
			4.4.1	4.5	5.1
PCI/ISA	2920	IBM PCI Token-Ring Adapter	X	X	X
	2931	ISA 8-Port Asynchronous Adapter	X	X	X
	2932	ISA 8-Port Asynchronous Adapter	X	X	X
	2933	ISA 128-Port Asynchronous Controller	X	X	X
	2741	PCI FDDI-Fiber Single-Ring Upgrade	X	X	X
	2742	PCI FDDI-Fiber Dual-Ring Upgrade	X	X	X
	2743	PCI FDDI-Fiber Single-Ring Upgrade	X	X	X
	2944	128-Port Asynchronous Controller, PCI bus	X	X	X
	2943	8-Port Asynchronous EIA-232/RS-422, PCI bus Adapter	X	X	X
	2963	Turboways 155 PCI UPT ATM Adapter	X	X	X
	2968	PCI Ethernet 10/100 Adapter	X	X	X
	2969	PCI Gigabit Ethernet Adapter	X	X	X
	2979	PCI AutoLANStreamer Token-Ring Adapter	X	X	X
	2985	PCI Ethernet BNC/RJ-45 Adapter	X	X	X
	2986	PCI Ethernet 10/100 Adapter	X	X	X
	2987	PCI Ethernet AUI/RJ-45 Adapter	X	X	X
	2988	Turboways 155 PCI MMF ATM Adapter	X	X	X
	4951	IBM 4-Port 10/100 Base-Tx Ethernet PCI Adapter	X	X	X
	4953	IBM 64-bit/66MHz PCI ATM 155 UTP Adapter		X	X
	4957	IBM 64-bit/66MHz PCI ATM 155 MMF Adapter		X	X
	4961	IBM Universal 4-Port 10/100 Ethernet Adapter		X	X
	4962	IBM 10/100 Mbps Ethernet PCI Adapter II		X	X
	4959	Token-Ring PCI Adapter	X	X	X

Type	Feature code	Description	HACMP version		
			4.4.1	4.5	5.1
PCI/ISA	8396	RS/6000 SP System Attachment Adapter	X	X	X
	2975	10/100/1000 Base-T Ethernet PCI Adapter		X	X
	5700	IBM Gigabit Ethernet-SX PCI-X Adapter			X
	5701	IBM 10/100 Base-TX Ethernet PCI-X Adapter			X
MCA	1904	Fibre Channel Adapter	X	X	X
	2402	Network Terminal Accelerator Adapter	X	X	
	2403	Network Terminal Accelerator Adapter	X	X	
	2723	FDDI-Fiber Dual-Ring Upgrade	X	X	X
	2724	FDDI-Fiber Single-Ring Adapter	X	X	X
	2725	FDDI-STP Single-Ring Adapter	X	X	X
	2726	FDDI-STP Dual-Ring Upgrade	X	X	X
	2930	8-Port Asynchronous Adapter - EIA-232	X	X	X
	2964	10/100 Mbps Ethernet Adapter - UNI	X	X	X
	2972	AutoLANStreamer Token-Ring Adapter	X	X	X
	2980	Ethernet High-Performance LAN Adapter	X	X	X
	2989	Turboways 155 ATM Adapter	X	X	X
	2992	Ethernet/FDX 10 Mbps TP/AUI MC Adapter	X	X	X
	2993	Ethernet BNC MC Adapter	X	X	X
	2994	10/100 Mbps Ethernet Adapter - SMP	X	X	X
	4018	High-Performance Switch (HPS) Adapter-2	X	X	X
	4020	Scalable POWERParallel Switch Adapter	X	X	X
	4025	Scalable POWERParallel Switch2 Adapter		X	X
	4025	Scalable POWERParallel Switch2 MX2 Adapter (#4026) on SP nodes (APARs IY25705, IY23767, and IY25322 required)		X	X
	2412	Enhanced SCSI-2 Differential Fast/Wide Adapter/A	X	X	X
2415	SCSI-2 Fast/Wide Adapter/A	X	X	X	

Type	Feature code	Description	HACMP version		
			4.4.1	4.5	5.1
MCA	2416	SCSI-2 Differential Fast/Wide Adapter/A	X	X	X
	2420	SCSI-2 Differential High-Performance External I/O Controller	X	X	X
	6212	High Performance Subsystem Adapter/A (40/80 Mbps)	X	X	X
	6214	SSA 4-Port Adapter	X	X	X
	6216	Enhanced SSA 4-Port Adapter	X	X	X
	6219	MCA SSA Adapter	X	X	X
PCI	6203	PCI Dual Channel Ultra3 SCSI Adapter		X	X
	6204	PCI Universal Differential Ultra SCSI Adapter	X	X	X
	6205	PCI Dual Channel Ultra2 SCSI Adapter	X	X	X
	6206	PCI SCSI-2 Single-Ended Ultra SCSI Adapter	X	X	X
	6207	PCI SCSI-2 Differential Ultra SCSI Adapter	X	X	X
	6208	PCI SCSI-2 Single-Ended Fast/Wide Adapter	X	X	X
	6209	PCI SCSI-2 Differential Fast/Wide Adapter	X	X	X
	6215	PCI SSA Adapter	X	X	X
	6225	Advanced SerialRAID Adapter	X	X	X
	6230	Advanced SerialRAID Plus Adapter (including Fast Write Cache (#6235) with two-initiator only)	X	X	X
	6227	Gigabit Fibre Channel Adapter	X	X	X
	6228	1- and 2-Gigabit Fibre Channel Adapter for 64-bit PCI Bus		X	X
	4025	Scalable POWERParallel Switch2 PCI Attachment Adapter (#8379) for SP attached servers (APARs IY25705, IY23767, and IY25322 required)		X	X

External storage subsystems

HACMP supports the external storage subsystems listed in Table 15-5.

Notes:

X Supported

Blank Not supported

Table 15-5 External storage subsystems

Subsystem	HACMP version		
	4.4.1	4.5	5.1
IBM 7131 SCSI Multi-Storage Tower Model 105 (supports up to four nodes; no CD-ROMs or tapes can be installed)	X	X	X
IBM 7131 SSA Multi-Storage Tower Model 405 (supports up to eight nodes; no CD-ROMs or tapes can be installed)	X	X	X
IBM 7133 SSA Disk Subsystem Models 020 and 600 (supports up to eight nodes)	X	X	X
IBM 7133 SSA Disk Subsystem Models D40 and T40 in up to 72.8 GB Mode (supports up to eight nodes)	X	X	X
IBM 7137 Disk Array Subsystem Models 413, 414, 415, 513, 514, and 515 (supports up to four nodes)	X	X	X
IBM 7204 External Disk Drive Models 317, 325, 339, 402, 404, and 418 (supports up to four nodes)	X	X	X
IBM 2105 Versatile Storage Server (VSS) Models B09 and 100 (supports up to four nodes)	X	X	X
IBM Enterprise Storage Server (ESS) Models E10, E20, F10, and F20 (supports up to eight nodes using SCSI and Fibre Channel interfaces using IBM FC/FICON Features 3021, 3022, and 3023)	X	X	X
IBM 2105 TotalStorage Enterprise Storage Server (ESS) Model 800			X
IBM TotalStorage FASiT200 Storage Server 3542-2RU			X
IBM TotalStorage FASiT500 Storage Server 3552-1RU			X
IBM TotalStorage FASiT700 Storage Server 1742-1RU			X
IBM TotalStorage FASiT900 Storage Server 1742-900			X
IBM 2102-F10 Fibre Channel RAID Storage Server	X	X	X
IBM 2104 Expandable Storage Plus Models DL1, TL1, DU3, and TU3	X	X	X

Subsystem	HACMP version		
	4.4.1	4.5	5.1
IBM 2108 Storage Area Network (SAN) Data Gateway Model G07	X	X	X
IBM 2031-016 McData ES-3016 Fibre Channel Product		X	X
IBM 2031-032 McData ES-3032 Fibre Channel Product		X	X
IBM 2031-232 McData ES-3232 Fibre Channel Product			X
IBM 2031-216 McData ES-3216 Fibre Channel Product			X
IBM 2032-001 McData ED-5000 Fibre Channel Director		X	X
IBM 2032-064 McData ED-6064 Fibre Channel Director		X	X
INRANGE FC/9000 Fibre Channel Director Model 2042-001 and Model 2042-128		X	X
IBM 2109 Models S08, S16, F16, F32 and M12 SAN Fibre Channel Switch		X	X
IBM 3534 Model F08 SAN Fibre Channel Switch			X
IBM 3583 Ultrium Scalable Tape Library model L18, L32 & L72		X	X
IBM 3584 Ultra Scalable Tape Library model L32 & D32			X
IBM TotalStorage Enterprise Tape Drive 3590 model H11			X
IBM Magstar 3590 Tape Drive model E11 & B11		X	X
IBM 3581 Ultrium Tape Autoloader model H17 & L17			X
IBM 3580 Ultrium Tape Drive model H11 & L11			X

Router support

HACMP supports the routers listed in Table 15-6 on page 660.

Notes:

X Supported

Blank Not supported

Table 15-6 Router support

Router support	HACMP version		
	4.4.1	4.5	5.1
IBM RS/6000 SP Switch Router 9077-04S	X	X	X
IBM 7139-111 Vicom Systems SLIC Router		X	X
IBM 7140-160 SAN Controller 160		X	X

Rack-mounted storage subsystems

Table 15-7 provides a list of rack-mounted storage subsystems that HACMP supports.

Notes:

X Supported
Blank Not supported

Table 15-7 Rack-mounted storage subsystems

Rack-mounted storage subsystems	HACMP version		
	4.4.1	4.5	5.1
IBM 7027 High Capacity Storage Drawer Model HSC (supports up to two nodes; no CD-ROMS or tapes installed)	X	X	X
IBM 7027 High Capacity Storage Drawer Model HSD (supports up to four nodes; no CD-ROMS or tapes installed)	X	X	X

15.1.12 Software requirements

The following sections provide an overview of the supporting relationship between HACMP and software.

HACMP Version 4.4.1

HACMP Version 4.4.1 requires the following:

- ▶ AIX
 - AIX Version 4.3.3 (or later modification levels) on pSeries or RS/6000 servers with at least four slots; PSSP Version 3.2 is also required on RS/6000 SP (9076) Systems with AIX Version 4.3, or
 - AIX 5L Version 5.1 (or later modification levels) on pSeries or RS/6000 (non-SP) servers with at least four slots, and APARs IY17684, IY19089, and IY19156.

► RSCT

If HACMP/ES Version 4.4.1 is to be installed on an AIX Version 4.3.3 system, the RSCT file sets delivered with HACMP must be installed (rsct.basic.hacmp 1.2.1.0 and rsct.clients.hacmp 1.2.1.0).

HACMP/ES 4.4.1 is also supported on AIX 5L Version 5.1. For AIX Version 5.1, the RSCT file sets will be delivered on the media with AIX. If HACMP/ES Version 4.4.1 is to be installed on an AIX 5L Version 5.1 system, the RSCT file sets delivered with AIX must be installed (rsct.basic.hacmp 2.2.0.0 and rsct.compat.clients.hacmp 2.2.0.0).

Each processor within a high-availability server complex requires the licensed program HACMP Version 4.4.1 to be installed. Except during the upgrade process from earlier releases of HACMP Version 4, all processors in the HACMP Version 4.4.1 server complex must be at the same AIX operating system level, including PTFs and maintenance upgrades.

Note: HACMP Version 4.4.1 has limitations on AIX 5L Version 5.1 as follows:

- Fencing for concurrent mode volume groups created on 9333 disks on AIX 5L Version 5.1 will not be supported.
- Enhanced concurrent volume groups created on AIX 5L Version 5.1 systems cannot be accessed on AIX Version 4.3, or earlier systems. They also cannot be accessed on nodes running HACMP Classic. This is because RSCT is a prerequisite for Enhanced Concurrent Volume Groups.
- HACMP can only run on 32-bit AIX kernels. Even if the hardware is capable of supporting 64-bit kernels, AIX needs to be initialized with a 32-bit kernel.
- Visual Systems Manager (VSM) will not be supported with AIX 5L Version 5.1.

HACMP Version 4.5

HACMP Version 4.5 requires the following:

► AIX

HACMP Version 4.5.0 requires AIX 5L Version 5.1 (or later modification levels), Version 5.2 on IBM @server Cluster 1600, pSeries, or RS/6000 servers with at least four slots, and APARs IY19089 and IY19156.

► RSCT

The RSCT filesets delivered with AIX must be installed (rsct.basic.hacmp 2.2.1.0 and rsct.compat.clients.hacmp 2.2.1.0) for ES and ESCRM.

► PSSP

PSSP Version 3.4 is also required on RS/6000 SP(9076) systems.

Each node within a high-availability server complex requires the licensed program HACMP Version 4.5.0 to be installed. Except during the upgrade process from earlier releases of HACMP Version 4, all nodes in the HACMP server complex must be at the same AIX operating system level, including PTFs and maintenance upgrades.

Some of the devices supported in HACMP Version 4.5.0 may require a later release level of the AIX operating system; refer to the specific hardware announcement for the AIX release levels required by the hardware.

HACMP Version 5.1

HACMP Version 5.1 requires the following:

► AIX

HACMP Version 5.1.0 requires AIX 5L on IBM @server pSeries, Cluster 1600, or RS/6000 servers with at least four slots.

The specific requirements for AIX 5L Version 5.1 or AIX 5L Version 5.2 are:

- AIX 5L Version 5.1 with the 5100-03 Recommended Maintenance package (or later modification levels)
- AIX 5L Version 5.2.0 with the 5200-01 Recommended Maintenance package 5.2.0.10 (or later modification levels)
- In addition, it is important to note that some HACMP utilities (clinfo and clstat) require SNMP Version 1 agents, even though the base AIX 5L Version 5.2 default is SNMPv3. See documentation APAR IY37779 for additional details.

► RSCT

The RSCT filesets delivered with AIX 5L must be installed. They are:

- AIX 5L Version 5.1: rsct.basic.hacmp 2.2.1.30 and rsct.compat.clients.hacmp 2.2.1.30
- AIX 5L Version 5.2: rsct.basic.hacmp 2.3.1.0 and rsct.compat.clients.hacmp 2.3.1.0

► PSSP

PSSP Version 3.5 is also required on RS/6000 SP (9076) systems.

- In order to use disk heartbeating for HACMP Version 5.1, customers will need these RSCT APARs:
 - RSCT on AIX 5L Version 5.1 IY42782

- RSCT on AIX 5L Version 5.2 IY42783

Each node within a high-availability server complex requires the licensed program HACMP Version 5.1.0 to be installed. Except during the upgrade process from earlier releases of HACMP Version 4, it is recommended that all nodes in the HACMP server complex be at the same AIX 5L operating system level, including PTFs and maintenance upgrades.

Some of the devices supported in HACMP Version 5.1.0 may require a later release level of the AIX 5L operating system; refer to the specific hardware announcement for the AIX 5L release levels required by the hardware.

- ▶ SDD (Subsystem Device Driver)
To use C-SPOC with VPATH disks, SDD Version 1.3.1.3 or later is required.
- ▶ HACMP/XD
HACMP/XD requires HACMP Version 5.1.0 (cluster.es.server.rte 5.1.0.0).
- ▶ ESS/PPRC mirroring
For ESS/PPRC mirroring:
 - AIX 5L Java Runtime Environment 1.3.0.15 or higher is required.
 - IBM ESS microcode level 2.1.1 or higher is required.
 - IBM 2105 Command Line Interface (ibm2105cli.rte 32.6.100.13) or the IBM 2105 Command Line Interface (ibm2105esscli.rte 2.1.0.15) is required.

Note: It is assumed that the command line interface is installed in its default location, /usr/opt/ibm2105cli.

- IBM 2105 Subsystem Device Driver (SDD) ibmSdd_510nchacmp.rte 1.3.3.6 or higher is required.
- ▶ For XD (HAGEO) IP-based mirroring, no additional prerequisites are required.

15.2 HAGEO/GeoRM

The IBM High Availability Geographic Cluster for AIX (HAGEO) and IBM Geographic Remote Mirror (GeoRM) software product provide a flexible, reliable platform for building disaster-tolerant computing environments. HAGEO components provide the capability for mirroring data across TCP/IP point-to-point networks over an unlimited distance from one geographic site to another. It works with the IBM High Availability Cluster Multi-Processing

(HACMP) licensed program product to provide automatic detection, notification, and recovery of an entire geographic site from failures. This section discusses some basic disaster recovery concepts and explains what the HAGEO/GeoRM software can do to help you prepare for disaster recovery and perform it.

15.2.1 Overview

HAGEO and GeoRM are real-time geographic clustering and data backup solutions providing mirroring of a customer's business-critical data over any distance. HAGEO integrates with the cluster management functions of IBM HACMP/ES to allow the building of systems that automatically detect site failures and provide automatic, timely recovery in the event of a physical disaster, such as an extended power outage, a fire, flood, hurricane, or earthquake.

GeoRM provides the same geo-mirroring functions as HAGEO, but does not provide automatic failover between sites. GeoRM allows up to eight IBM @server pSeries and RS/6000 computers to be placed in widely separated geographic locations.

Designed to operate over unlimited distances, the HAGEO Version 2.4 and GeoRM Version 2.4 data replication environments extend to worldwide areas through remote data mirroring (geo-mirroring) over any IP-based communications paths, including Internet, token ring, Ethernet, FDDI, T1, and ATM.

Table 15-8 provides a brief history of HAGEO.

Table 15-8 HAGEO releases and features

HAGEO release	Enhancements
2.1	<ul style="list-style-type: none"> ▶ First IBM “badged” release of HAGEO. ▶ Synchronous, asynchronous, and MWC mirroring. ▶ Support for cascading, rotating, and concurrent resource groups. ▶ Several configuration restrictions. ▶ HACMP “Classic” support only.

HAGEO release	Enhancements
2.2	<ul style="list-style-type: none"> ▶ First IBM owned release of HAGEO. ▶ Geographic Remote Mirroring (GeoRM) introduced. GeoRM is a disaster recovery software that includes point-to-point mirroring of critical data without the automatic takeover of applications on the takeover site. ▶ Configuration from a single point. ▶ Tunable parameters added. ▶ Improved tracing and logging. ▶ NLS support. ▶ Support for HACMP ES.
2.3	<ul style="list-style-type: none"> ▶ Performance and management improvements. ▶ Support for large GeoMirror Device (GMD) sizes. ▶ Ability to resize geo-mirrors dynamically. ▶ Simpler and faster configuration. ▶ GMD state is preserved during resynchronization. ▶ Local peers in asynchronous mode. ▶ HAGEO operations log to HACMP log files. ▶ More flexible data transmission rates.
2.4	<ul style="list-style-type: none"> ▶ Latest release of HAGEO. ▶ Integration of HAGEO with HACMP ES. ▶ Dropped support for HACMP HAS (Classic). ▶ GeoManager removed from HAGEO. ▶ Option to use TCP instead of UDP protocol for mirroring. ▶ 64-bit kernel support.

15.2.2 Hardware requirements

The following sections describe the hardware requirements for HAGEO and related software.

HAGEO

HAGEO Version 2.4, with HACMP/ES Version 4.5, runs on IBM @server pSeries and RS/6000 uniprocessors, SMP servers, and Scalable POWERparallel® (SP) systems in a no single-point-of-failure server configuration. HAGEO Version 2.4 supports the IBM @server pSeries, Cluster 1600, and RS/6000 models designed for server applications that meet the AIX 5L and HACMP Version 4.5 requirements for internal memory, internal disk, and I/O

slots. HAGEO Version 2.4 is supported on all pSeries models, RS/6000 models, and devices supported by HACMP (5765-E54) Version 4.5.

The minimum configuration and sizing of each machine is highly dependent on the user's application and performance requirements. In configuring an HAGEO system, give particular attention to the following:

- ▶ Fixed disk capacity and mirroring (Logical Volume Manager and database)
- ▶ Other LAN devices (routers, bridges) and their effect on the system environment
- ▶ Other network software

HAGEO requires an IP-based communications network. HAGEO Version 2.4 installation requires 5 MB of disk storage space.

GeoRM

GeoRM Version 2.4 executes with IBM @server Cluster 1600 pSeries and RS/6000 uniprocessors, SMP servers, and Scalable POWERparallel (SP) systems. It supports pSeries and RS/6000 models designed for server applications that meet the AIX 5L minimum requirements for internal memory, internal disk, and I/O slots.

GeoRM Version 2.4 is supported on all pSeries models, RS/6000 models, and devices supported by AIX 5L Version 5.1 or later. Any pSeries and RS/6000 or device supported on AIX 5L Version 5.1 or later can be joined with any other supported RS/6000 or device in a GeoRM Version 2.4 configuration.

The minimum configuration and sizing of each machine is highly dependent on the user's application and performance requirements. In configuring a GeoRM system, particular attention must be paid to:

- ▶ Fixed-disk capacity and mirroring (Logical Volume Manager (LVM) and database)
- ▶ Other LAN devices (routers and bridges) and their effect on the system environment
- ▶ Other network software

GeoRM requires an IP-based communications network. GeoRM Version 2.4 installation requires 4 MB of disk storage space.

15.2.3 Software requirements

The following sections discuss the software requirements of HAGEO and related software.

HAGEO

HAGEO Version 2.4 requires AIX 5L Version 5.1 (5765-E61) or AIX 5L Version 5.2 (5765-E62) and HACMP ES Version 4.5.0 or later, and APARs IY34727, IY34728, IY34662, IY34845, and IY34846. Each processor within a high-availability geographic cluster requires these licensed programs to be installed.

Refer to the announcements for AIX 5L Version 5.1 and HACMP Version 4.5 for requirements regarding those products.

GeoRM

GeoRM Version 2.4 requires AIX 5L Version 5.1 (5765-E61) or AIX 5L Version 5.2 (5765-E62) or later, and APARs IY34727, IY34662, IY34845, and IY34846. GeoRM Version 2.4 should neither be installed nor run on the same system as HAGEO. Each server requires the licensed program to be installed. Refer to the individual software announcements for AIX 5L Version 5.1 or later for detailed requirements regarding that product.

15.2.4 Configuration examples

We will discuss the three most popular HAGEO designs in the following sections:

- ▶ Two nodes at each site
- ▶ Two nodes at the primary site and one node at the backup site
- ▶ One node at each site

Two nodes at each site

This design is the most recommended one and gives the most flexibility. The primary and also the backup site has a full HACMP cluster installed. All component failures are handled locally at each site. Only site failures result in failover across sites. Figure 15-5 on page 668 gives an overview of such a setup.




Figure 15-5 Two nodes at each site

Two nodes at the primary and one node at the backup

This is not a preferred design, but it is more popular. We handle all component failures locally only at the primary site. The secondary site propagates any node failures to a site failure. Even though the backup site can be configured to fail over anything from the primary, this design is more satisfactory if there are no applications running at the backup site, which avoids unnecessary disruption in case of maintenance on the backup. Figure 15-6 on page 669 gives an overview of this setup.

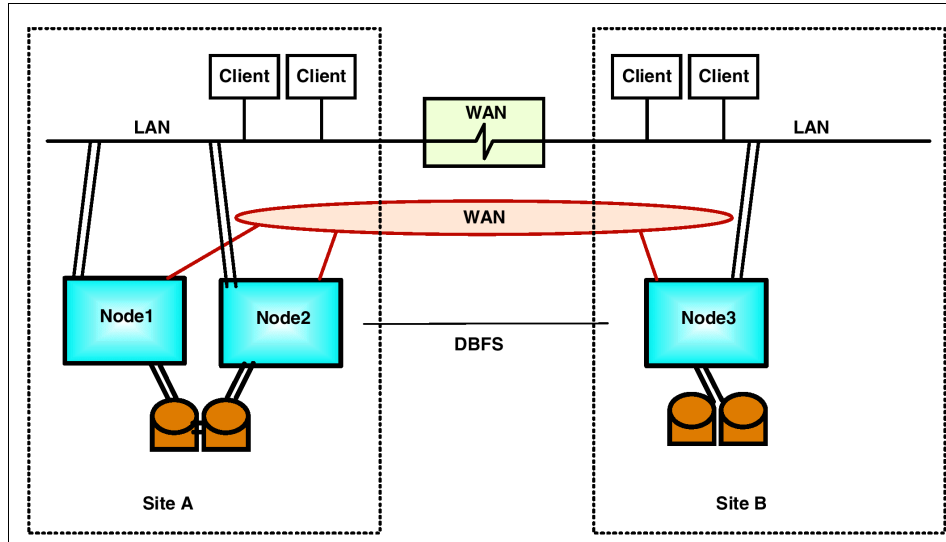


Figure 15-6 Two nodes at the primary site and one node at the backup site

One node at each site

This design is a very restricted one. It may handle disk or adapter failures, but node failures are propagated to site failures, because there are no local peer nodes available in this model. Figure 15-7 gives an overview of this setup.




Figure 15-7 One node at each site

15.2.5 New features of HAGEO Version 2.4/GeoRM Version 2.4

The new version of HAGEO/GeoRM provides world-class disaster recovery and data mirroring capabilities spanning unlimited distance, and automatic switch-over and resynchronization ensure uninterrupted access to data and applications.

HAGEO Version 2.4 and GeoRM Version 2.4 enhancements are as follows:

- ▶ Improved network performance capability with the addition of the TCP transport option, which is a new choice for the IP subsystem. UDP, which was the only choice available in previous releases, continues to be supported. UDP yields the best performance when there is a single LAN segment between sites. TCP provides the best performance when one or more routers or bridges are in the path between sites.
- ▶ The TCP transport option includes a new additional choice in the temporal ordering of mirrored data: write order by volume groups. User-configurable write ordering choices include:
 - Complete ordering by system (system-wide)
 - Write ordering by volume group (only with TCP)
 - None (only with TCP)

Writes across multiple GeoMirror devices can be mirrored on the remote node in the same order as the application wrote the data, as in previous releases, or not. The decision on whether to require system-wide ordering or order by volume group depends on the way data is organized and related. When ordering is not required, a performance gain may be realized

- ▶ 64-bit kernel support for the TCP transport option. In order to exploit this support, however, you need to migrate all of your nodes to HAGEO Version 2.4, and you need to select TCP as your communications transport option.

HAGEO Version 2.4 unique enhancements are:

- ▶ An automatic use of HACMP Configuration Data. HAGEO Version 2.4 increases focus on ease of use with tighter integration with IBM High Availability Cluster Multi-Processing (HACMP 5765-E54). Exploiting the HACMP infrastructure, HAGEO can take advantage of the HACMP cluster site configuration definitions. HAGEO Version 2.4 now uses your existing HACMP Configuration Data to:
 - Simplify the definition of the cluster topology and resource groups
 - Eliminate redundant data collection and data entry errors
 - Reduce configuration setup time
 - Ensure HACMP/HAGEO consistency

- Reduce cluster administration costs

HAGEO and GeoRM have three functional components (each):

- ▶ Geographic Mirroring (GeoMirror) works on top of the AIX logical volume disk manager. It replicates data, regardless of the applications involved, and coordinates updates at each location, ensuring that sites are synchronized.
- ▶ Geographic Messaging (GeoMessage) provides reliable communications between sites.
- ▶ Geographic Cluster Manager (GeoManager) differs in HAGEO and GeoRM. GeoManager for HAGEO acts as a supervisor for the distributed cluster: it reacts to network and site-related events, such as a site failure or site isolation. HAGEO is designed for installation with HACMP and is the GeoManager component that integrates with that program, allowing automatic detection and response to significant events in the geographic cluster without customer assistance. GeoManager for GeoRM provides tools to allow the user to manually manage the failover and reintegration of the configuration following a site or machine failure.

Also, a geo-mirror operates in one of three possible modes:

- ▶ Synchronous
- ▶ Synchronous with Mirror Write Consistency
- ▶ Asynchronous

In Synchronous mode, data is written first to the remote site and then at the local site where the transaction was 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 Mirror Write Consistency (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 helps to ensure 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 with a resulting longer period of time when the data on the local disk and the remote mirror differ. In the event of a site failure, some transactions may be lost, as most DB replicators are asynchronous.

Ultimately, customers should decide what mode is appropriate for their business, based on the trade-off between exact mirroring of data versus application performance. You are not forced to protect all your data in one way; you are allowed flexible performance and availability configurations to support your specific needs.

15.3 Cluster 1600

The IBM @server Cluster 1600 system delivers solutions to some of the most complex and largest technical, commercial applications by simultaneously bringing dozens of RISC processing nodes to a computing problem. This parallel processing capability enhances computing performance and throughput many times over serial computing.

In this section, we will discuss general concepts and enhancements of Cluster 1600.

15.3.1 Overview

The 9078 IBM @server Cluster 1600 is the official name for pSeries and/or RS/6000 clusters running AIX and builds on the success of the RS/6000 SP in high performance computing, large-scale database applications, and workload/server consolidation. The Cluster 1600 brings together cluster-ready servers or nodes, CSM or PSSP cluster management software, and a choice of IBM or industry standard cluster interconnect technologies into a unified cluster with its own unique identification number.

As cluster management software components, PSSP Version 3.5 on AIX 5L Version 5.1 and 5.2 for existing or new PSSP, High Performance Computing (HPC), and commercial customers, and Cluster Systems Management (CSM) on AIX 5L Version 5.2 and AIX 5L Version 5.1 (Maintenance Level 3 or later) for new customers are offered. Both share the same heritage, the IBM RS/6000 SP supercomputer, and both allow clusters of up to 128 servers (or compute nodes), or larger by special bid, to be controlled.

Also, specialized software is required for running, managing, and scheduling parallel supercomputing code and jobs. Cluster 1600 offers Parallel Environment Version 3.1 for writing and controlling parallel codes, LoadLeveler Version 3.1 for scheduling jobs, and ESSL Version 3.3 and Parallel ESSL Version 2.3 libraries for scientific and engineering calculations. All these are supported by the PSSP Version 3.5 management functions.

Cluster 1600 supports two generations of high-speed interconnect switch hardware to connect cluster nodes together: the SP Switch and the SP Switch2.

Although primarily aimed at the HPC customer, these interconnects support IP and thus can be used as a high-bandwidth LAN for purposes of backups, database connectivity, or other bandwidth-hungry applications. In addition, some databases, such as the parallel implementation of DB2, are designed to run a query across multiple nodes in a cluster simultaneously, with each node accessing a subset of the relational tables required for the query. These databases can make use of the switch interconnect for fast communication of query results back to the coordinating node. Data warehouses with dozens of terabytes of online data have been implemented in this fashion.

Cluster file system

GPFS for AIX Version 2.1 allows multiple systems to concurrently access a large file system over multiple I/O paths, while preserving standard UNIX file system semantics.

The Cluster 1600 is ideal for enterprise data centers that want to consolidate servers and workloads in order to increase management efficiency and lower cost. It is also ideal for organizations that want to scale workloads beyond the capacity of the largest available symmetric multiprocessor (SMP) servers. These include large-scale databases in commercial enterprises or very large computational models in scientific and technical computing environments.

15.3.2 Parallel System Support Program Version 3.5

Parallel System Support Program (PSSP) Version 3.5 offers the following enhancements and changes from previous releases:

- ▶ The 64-bit kernel version of AIX 5L Version 5.1, including any 64-bit kernel extensions and device drivers, can be used within a PSSP Version 3.5 cluster. The entire software stack supported by PSSP Version 3.5, GPFS, MPI, LoadLeveler, Parallel ESSL, LAPI, KLAPI, and VSD works with AIX 5L Version 5.1 or later 64-bit kernels, device drivers, and kernel extensions.
- ▶ Virtual Shared Disk (VSD) has been enhanced for performance in IP-based mode and adds improved diagnostics.
- ▶ Switch support has been enhanced to allow an administrator to specify sets of nodes to be excluded from serving as switch primary or primary backup.
- ▶ IBM now supports PSSP Version 3.5 on AIX 5L Version 5.2 (AIX 5L Version 5.2 support will not be offered in PSSP Version 3.4.) Also, PSSP Version 3.5 does not support AIX Version 4.3.3.
- ▶ PSSP Version 3.5 (and PSSP Version 3.4) will be able to utilize redundant HMCs attached to a logically partitionable server, providing better availability.

Note: PSSP Version 3.5 will be the last release of PSSP. For more information about the PSSP and CSM plans, contact your IBM technical representative, or see:

<http://www.ibm.com/servers/eserver/clusters/software/>

15.3.3 Cluster Systems Management Version 1.3.1

Cluster Systems Management (CSM) Version 1.3.1 is available for Cluster 1600 customers that are more interested in manageability clusters than performance clusters. Customers that do not have a cluster managed by PSSP or need a switch interconnect, and yet desire the cluster system management capability that PSSP customers have known, can choose CSM. CSM requires AIX 5L Version 5.2 for the management server and AIX 5L Version 5.1 with Maintenance Level 3 or later for the managed nodes.

15.3.4 General Parallel File System for AIX Version 2.1

The following features are available in General Parallel File System (GPFS) Version 2.1:

- ▶ The AIX 64-bit kernel is supported and exploited.
- ▶ Customers using GPFS Version 2.1 and PSSP Version 3.4 or 3.5 and AIX 5L Version 5.1 no longer need to purchase and install HACMP as a prerequisite. GPFS Version 2.1 can utilize the RSCT technology in AIX 5L Version 5.1 instead of HACMP.

15.3.5 Models and features

The Cluster 1600 provides a unique serial number under which all of the servers in this cluster will share. The servers within the 9078-160 serial number will keep their existing machine type/models. The 9078-160 acts as a second machine type/model.

A single system number that the Cluster 1600 has simplifies cluster administration and makes it easier for customers to track, manage and upgrade their clusters over time. Existing PSSP-managed clusters, including RS/6000 SP systems, can be incorporated into new Cluster 1600s, thus providing investment protection for existing cluster customers.

Features of this new model are content features that represent the servers, switches, and control workstation that comprise the cluster.

Table 15-9 and Table 15-10 show the announced features of the IBM @server Cluster 1600 (9078-160).

Table 15-9 Features of Cluster 1600

Description	Machine type	Model	Feature
7017 type server	9078	160	0001
7026 type server	9078	160	0002
9076-555/557 type (SPS)	9078	160	0003
9076-556/558 type (SPS2)	9078	160	0004
9076-SP	9078	160	0005
9076 SP expansion frame	9078	160	0006
Control Workstation	9078	160	0007
7040 type server	9078	160	0008
LPAR 7040 (switched)	9078	160	0009
7039 Type Server	9078	160	0010
LPAR 7039 (switched)	9078	160	0011
7028 type server	9078	160	0012
LPAR 7028 (switched)	9078	160	0013
7038 Type Server	9078	160	0014
LPAR 7038 (switched)	9078	160	0015

Table 15-10 Description of feature codes

Feature code	Description
0001	Includes the 7017-S85, 7017-S80, and 7017-S7a.
0002	Includes the 7026-6M1, 7026-6H1, 7026-6H0, 7026-H80, and 7026-M80.
0003	Includes the 9076-555 and the 9076-557 SP Switch Models.
0004	Includes the 9076-556 and the 9076-558 SP Switch2 Models. The SP Switch2 also supports two-plane switch fabrics with these models.

Feature code	Description
0005	Includes the 9076-550 and the 9076 Legacy Model Frame.
0006	Includes expansion frames on the 9076-550, -556, or Legacy Models.
0007	Includes the Control Workstation including the 7025-6F1, the 7026-6H1, the 7028-6E1, the 7028-6C1, the 7028-6C4 and the 7044-170.
0008	Includes the 7040-681 or 7040-671.
0009	Includes a switched LPAR of the 7040-681 or 7040-671.
0010	Includes the 7039-651.
0011	Includes a switched LPAR of the 7039-651.
0012	Includes the 7028-6C4.
0013	Includes a switched LPAR of the 7028-6C4.
0014	Includes the 7038-6M2.
0015	Includes a switched LPAR of the 7038-6M2.

15.3.6 Cluster enhancements

New offerings for the Cluster 1600 include additional hardware models as well as support for Parallel Systems Support Programs (PSSP) Version 3.5 on AIX 5L Version 5.2.

► pSeries 670 and 690 servers

Support on these servers with RIO-2 is available with PSSP Version 3.4 and PSSP Version 3.5. I/O performance is enhanced with the PCI-X and RIO-2 technologies. New, faster processors based on IBM POWER4+ technology offer larger memory and I/O enhancements for the pSeries 670 and 690 servers. PSSP Version 3.4, PSSP Version 3.5, and CSM Version 1.3.1 are supported on these servers, providing increased server performance and capacity.

► pSeries 655 servers

These servers are supported with PSSP Version 3.4 and Version 3.5 running AIX 5L Version 5.1 and PSSP Version 3.5 running AIX 5L Version 5.2 with the SP Switch2 in both single and two-plane environments. The servers are also supported with CSM Version 1.3.1 running AIX 5L Version 5.2. PSSP Version 3.4, PSSP Version 3.5, and CSM Version 1.3.1 support the pSeries 655 POWER4+ server with the RIO-2.

- ▶ pSeries 650 Cluster support

The pSeries 650 server is initially supported as a cluster building block with the Cluster 1600 and PSSP Version 3.5 on AIX 5L Version 5.2 providing LPAR support in a clustered environment. With PSSP, this server can optionally be attached to the SP Switch2. The pSeries 650 is to be supported with PSSP Version 3.4 and Version 3.5 on AIX 5L Version 5.1, and CSM Version 1.3.1.
- ▶ pSeries 630 Server

LPAR support in a clustered environment is available on the pSeries 630 server running either CSM Version 1.3.1, PSSP Version 3.4, or PSSP Version 3.5.
- ▶ Dual Hardware Management Console (HMC) support

CSM for AIX Version 1.3.1 provides dual HMC support, including power on/off LPAR control, booting, starting, and stopping the system or individual partitions, and displays the list of partitions and the power status.
- ▶ CSM hardware control features

CSM for AIX Version 1.3.1 further enhances hardware control support to include:

 - 9076 SP nodes (FC 2050, 2051, 2052, 2053, 2054, 2056, 2057, and 2058)
 - 7026 servers (Models H80, M80, 6H0/6H1, 6M1, and p660)
 - pSeries 690, 670, 655, 650, and 630 servers

Support on the SP node (FC 2054 and 2058), is planned to be available on September 26, 2003, with APAR IY42847.

Note: CSM does not support the SP Switch or SP Switch2.

- ▶ CSM interoperability

CSM for AIX 5L Version 1.3 supports interoperability with the AIX management server and managed xSeries server nodes running:

 - Red Hat Linux Version 7.2, 7.3, or 8.0
 - Red Hat Advanced Server Version 2.1
 - SuSE Linux Version 8.0 or 8.1
 - SuSE Linux Enterprise Server (SLES) Version 7 or 8

AIX 5L Version 5.2 is required for the cluster management server. The maximum number of operating system images (AIX plus Linux) that can be managed under CSM for AIX 5L is 128.

Table 15-11 shows currently available hardware products in terms of PSSP and CSM.

Table 15-11 Support availability

Program name and version		Supported H/W
PSSP	3.4	<ul style="list-style-type: none"> ▶ On pSeries 690 with RIO-2 ▶ On pSeries 630 with the SP switch2 PCI-X adapter with RIO ▶ On pSeries 670 POWER4+ and pSeries 690 POWER4+ ▶ On the pSeries 650 with RIO-2 and the SP Switch2 PCI-X Attachment adapter ▶ On the pSeries 655 POWER4+
	3.5	<ul style="list-style-type: none"> ▶ On AIX 5L Version 5.2 with pSeries 690, 670, 655, 650, and 630 servers ▶ On pSeries 690 with RIO-2 ▶ On AIX 5L Version 5.2 with pSeries 630 with the SP Switch2 PCI-X adapter with RIO ▶ On pSeries 670 POWER4+ and pSeries 690 POWER4+ ▶ On the pSeries 650 with RIO-2 and the SP Switch2 PCI-X Attachment adapter ▶ On the pSeries 655 POWER4+
CSM	1.3.1	<ul style="list-style-type: none"> ▶ On the pSeries 650 and 655 ▶ On the pSeries 670 POWER4+ and pSeries 690 POWER4+ ▶ On the pSeries 655 and 615 POWER4+ ▶ On the 9076 SP nodes and 7026 servers

15.3.7 Cluster 1600 with PSSP

The PSSP distributed server management technology has been used for commercial and high-performance computing environments for server and workload consolidation through a single-point-of-control for years. PSSP Version 3.5 is intended for existing SP and Cluster 1600 customers using PSSP, as well as new High Performance Computing (HPC) cluster customers, while CSM is intended for new Cluster 1600 customers in the commercial and HPC space. PSSP currently supports the software that High Performance Computing (HPC)

customers require. This includes GPFS, LoadLeveler, Parallel Environment, Parallel ESSL, and ESSL.

Expansibility of clusters

A Cluster 1600 with PSSP can consist of two to 128 AIX operating system images. An operating system image, or logical node, can be:

- ▶ 7040 server (p670, p690) running as a full system partition
- ▶ LPAR of a 7040 server (p670, p690)
- ▶ 7039 server (p655) running as a full system partition
- ▶ LPAR of a 7039 server (p655)
- ▶ 7038 server (p650)
- ▶ LPAR of a 7038 server (p650)
- ▶ 7028 server (p630)
- ▶ LPAR of a 7028 server (p630)
- ▶ 7026 server (H80, M80, p660, 6H0/6H1, 6M1)
- ▶ 7017 server (S70, S7A, S80, p680)
- ▶ 9076 SP node

Logical nodes are limited in a cluster running PSSP, as described in Table 15-12 and the scaling rules that follow.

Table 15-12 Cluster limits for PSSP

Maximum values	SP(1) Switch	SP(1) Switch 2	SP Switch 2 (2)	Industry standard inter connect
p690/p670 servers per cluster	32	32	32	32
p655 servers per cluster	-	64	64	64
p650 servers per cluster	-	64	64	64
p630 servers per cluster	-	64	64	64
LPARs per p690 server	8	16	16	16
LPARs per p670 server	4	4	16	16
LPARs per p655 server	-	2	4	4
LPARs per p650 server	-	2	8	8
LPARs per p630 server	-	2	4	4
LPARs per cluster	128	128	128	128

Maximum values	SP(1) Switch	SP(1) Switch 2	SP Switch 2 (2)	Industry standard inter connect
Number of switch planes supported				
p690/p670	1	1 or 2	1 or 2	0
p655	-	1 or 2	1 or 2	0
p650	-	1	1	0
p630	-	1	1	0
Number of Servers per HMC (3)				
p690/p670	8	8	8	8
p655(4)	-	16	16	16
p650(4)	-	16	16	16
p630(4)	-	16	16	16
Number of LPARs per HMC	32	32	32	32

Notes:

1. SP Switch2 supports pSeries 690/pSeries 670 servers with both Power4 and Power4+ MCMs. SP Switch support for pSeries 690/pSeries 670 servers is limited to Power4 MCMs.
2. With the SP Switch, switched and non-switched logical nodes cannot be mixed.
3. Redundant HMCs are supported.
4. The HMC can control servers of multiple machine types. When combining multiple machine types on an HMC, substitutions of other server types are based upon the relative weighting factor of the server. For example, a pSeries 690 is twice the weighted factor of a pSeries 630 server. A supported HMC configuration could control four pSeries 690 servers and eight pSeries 630 servers.

Prerequisites

A cluster 1600 with PSSP must meet all of the following prerequisites:

- ▶ No more than 128 logical nodes from the set (7040, 7039, 7038, 7028, 7026, 7017, and 9076)
- ▶ No more than 32 servers or 128 LPARs from the set (7040)

- ▶ No more than 64 servers from the set (7040, 7039, 7038, 7028, 7026, and 7017)
- ▶ No more than 16 servers from the set (7017)
- ▶ No more than 128 9076 SP nodes

15.3.8 Cluster 1600 with CSM

CSM Version 1.3.1 for AIX is available and shipping with AIX 5L Version 5.2. When customers consider using CSM instead of PSSP for building their Cluster 1600, then customers need to understand the differences between the products until CSM provides the same functionality as PSSP. CSM is supplied try-and-buy for those who are new to clustering and want to try it out. As new customers will probably want to make use of the new features of AIX 5L Version 5.2, such as dynamic logical partitioning on the pSeries POWER4 machines, CSM is the better choice if the cluster is needed now.

Expansibility of clusters

A Cluster 1600 with CSM can consist of two to 128 AIX operating system images. An operating system image, or logical node, can be:

- ▶ 7040 server (p670, p690) running as a full system partition
- ▶ LPAR of a 7040 server (p670, p690)
- ▶ 7039 server (p655) running as a full system partition
- ▶ LPAR of a 7039 server (p655)
- ▶ 7038 server (p650)
- ▶ LPAR of a 7038 server (p650)
- ▶ 7029 server (p615 Model 6C3 only)
- ▶ 7028 server (p630 Model 6C4 only)
- ▶ LPAR of a 7028 server (p630 Model 6C4 only)
- ▶ 7026 server (H80, M80, p660, 6H0/6H1, and 6M1)
- ▶ 9076 SP node (9076 machine type, feature numbers 2050, 2051, 2052, 2053, 2054, 2056, 2057, and 2058)

Logical nodes are limited in a cluster running CSM, as described in Table 15-13 on page 682 and the scaling rules that follow.

Table 15-13 Cluster limits for CSM

Maximum values	Non-switched	Industry standard inter connect
p690/p670 servers per cluster	32	32
p655 servers per cluster	64	64
p650 servers per cluster	64	64
p630 servers per cluster	64	64
p615 servers per cluster	64	64
LPARs per p690 server	16	16
LPARs per p670 server	16	16
LPARs per p655 server	4	4
LPARs per p650 server	8	8
LPARs per p630 server	4	4
LPARs per cluster	128	128
Number of Servers per HMC (1)		
p690/p670	8	8
p655(2)	16	16
p650(2)	16	16
p630(2)	16	16
p615(2)	16	16
Number of LPARs per HMC	16	16

Notes:

1. Redundant HMCs are supported.
2. The HMC can control servers of multiple machine types. When combining multiple machine types on an HMC, substitutions of other server types are based upon the relative weighting factor of the server. For example, a pSeries 690 is twice the weighted factor of a pSeries 630 server. A supported HMC configuration could control four pSeries 690 servers and eight pSeries 630 servers.

Prerequisites

A cluster 1600 with CSM must meet all of the following prerequisites:

- ▶ No more than 128 logical nodes from the set (7040, 7039, 7038, 7028, 7026, and 9076)
- ▶ No more than 32 servers or 128 LPARs from the set (7040)
- ▶ No more than 64 servers from the set (7040, 7039, 7038, 7029, and 7028)



Resource management

This chapter provides an overview of resource management for pSeries, with the following sections:

- ▶ Partitioning including LPAR and DLPAR
- ▶ IBM Hardware Management Console for pSeries (HMC)
- ▶ Service Agent
- ▶ Service Focal Point
- ▶ Inventory Scout

16.1 Partitioning

In this section, we introduce necessary concepts and terminology to understand the logical partitioning implementation on the IBM @server pSeries servers. We also explain the logical partitioning implementation on partitioning-capable pSeries servers and dynamic logical partitioning. For more information, see *The Complete Partitioning Guide for IBM @server pSeries Servers*, SG24-7039.

This section contains the following sections:

- ▶ Logical partitioning overview
- ▶ Partitioning implementation on pSeries servers
- ▶ Dynamic Logical Partitioning
- ▶ Managing partitions

16.1.1 Logical partitioning overview

In this section, we first introduce necessary concepts and terminology to understand the logical partitioning implementation on the IBM @server partitioning-capable pSeries servers. We discuss the following topics:

- ▶ Partitioning implementations
- ▶ Partitioning support on pSeries servers
- ▶ LPAR configuration notes
- ▶ Terminology used in partitioning
- ▶ Four terms regarding memory

16.1.2 Partitioning implementations

There is a strong demand for high-end systems to provide greater flexibility, in particular, the ability to subdivide them into smaller partitions that are capable of running a version of an operating system or a specific set of application workloads.

IBM initially started work on partitioning S/370™ mainframe systems in the 1970s. Since then, logical partitioning (LPAR) on IBM mainframes (now called IBM @server zSeries) has evolved from a predominantly physical partitioning scheme, based on hardware boundaries, to one that allows for virtual and shared resources with dynamic load balancing. In 1999, IBM implemented LPAR support on the AS/400 (now called IBM @server iSeries) platform. In 2000, IBM announced the ability to run the Linux operating system in an LPAR on a zSeries server.

Throughout this redbook, we refer to the different partitioning mechanisms available on the market. Therefore, it is appropriate to clarify the following terms and definitions by which we classify these mechanisms:

- Building block** A collection of system resources, such as CPUs, memory, and I/O connections. These may be physically packaged as a self-contained symmetric multiprocessing (SMP) system (rack-mounted or stand-alone) or as boards within a larger multiprocessor system. There is no requirement for the CPUs, memory, and I/O slots to occupy the same physical board within the system, although they often do. Other vendors use the terms system board, cell, and Quad Building Block (QBB) to refer to their building blocks.
- Physical partition** One or more building blocks linked together by a high-speed interconnect. Generally, the interconnect is used to form a single, coherent memory address space. In a system that is only capable of physical partitioning, a partition is a group of one or more building blocks configured to support an operating system image. Other vendors may refer to physical partitions as *domains* or *nPartitions*.
- Logical partition** A subset of logical resources that are capable of supporting an operating system. A logical partition consists of CPUs, memory, and I/O slots that are a subset of the pool of available resources within a system.

Note: The major difference between logical partitioning and physical partitioning is the granularity and flexibility available for allocating resource to an operating system image. Logical partitions have finer granularities than physical partitions.

It should be noted that the zSeries LPAR implementation is unique in comparison to the other partitioning implementations available from IBM and other hardware vendors. It is a mature and dynamic partitioning technology. The experience of IBM with physical and logical partitioning over the last 25 years has greatly influenced the design and implementation of logical partitioning on pSeries.

The pSeries 690 server is the first pSeries server to incorporate the ability of being partitioned. Its architectural design brings logical partitioning to the UNIX world, being capable of multiple partitions inside a single server, with great flexibility in resource selection. The partitioning implementation on pSeries 690 differs from those of other UNIX system vendors in that the physical resources that can be assigned to a partition are not limited by internal physical system

board boundaries. Now, IBM is expanding its partitioning-capable pSeries server lineup, as explained in 16.1.3, “Partitioning support on pSeries servers” on page 688.

Processors, memory, and I/O slots can be allocated to any partition, regardless of their locality. For example, two processors on the same POWER4 silicon chip can be in different partitions. Peripheral component interconnect (PCI) slots are assigned individually to partitions, and memory can be allocated in fixed-size increments. The fine granularity of the resources that can be assigned to partitions provides flexibility to create systems by the desired resources.

The partitioning-capable pSeries servers are also capable of running both AIX 5L Version 5.1 or later and Linux inside a partition on the single system simultaneously.

Many of the features described in this document are operating system dependant and may not be available on Linux. For more information, see:

http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html

Note: Hereafter, we refer to logical partitions as partitions throughout this redbook.

16.1.3 Partitioning support on pSeries servers

In addition to the first partitioning-capable pSeries server, pSeries 690, IBM now is expanding the partitioning-capable pSeries server lineup, as explained in this section.

Supported models

At the time of writing, the IBM @server pSeries servers shown in Table 16-1 support partitioning.

Table 16-1 Supported partitioning-capable pSeries servers

Official product model name	Short product name	M/T-MDL
IBM @server pSeries 690 Model 681	pSeries 690	7040-681
IBM @server pSeries 670 Model 671	pSeries 670	7040-671
IBM @server pSeries 655 Model 651	pSeries 655 Model 651	7039-651

Official product model name	Short product name	M/T-MDL
IBM @server pSeries 650 Model 6M2	pSeries 650 Model 6M2	7038-6M2
IBM @server pSeries 630 Model 6C4	pSeries 630 Model 6C4	7028-6C4
IBM @server pSeries 630 Model 6E4	pSeries 630 Model 6E4	7028-6E4

Note: The IBM @server 7029 pSeries 615 can be attached to a HMC as well, but the system does not support partitioning.

Depending on supported number of processors, the maximum number of partitions are shown in Table 16-2.

Table 16-2 Maximum number of processors, memory size, and partitions

Short product name	Maximum number of processors	Maximum memory size in GB	Maximum number of I/O drawers	Maximum number of partitions
pSeries 690	32 ^a	512	8	16
pSeries 670	16	128	3	16
pSeries 655 Model 651	8	32	1	2
pSeries 650 Model 6M2	8	64	8	8 ^b
pSeries 630 Model 6C4	4	32	2	4 ^c
pSeries 630 Model 6E4	4	32	0	2

a. The High Performance Computing (HPC) feature of pSeries 690 is equipped with up to 16 processors.

b. Needs external disk subsystems for the boot disk.

c. When equipped with I/O drawers.

The logical partitioning concept and required tasks are basically similar on these partitioning-capable pSeries server models. However, there is a substantial difference when assigning I/O resources to partitions depending on the models. For the hardware model-specific information about the I/O resource assignment, see 16.2.3, “I/O device assignment considerations” on page 717.

16.1.4 LPAR configuration notes

In this section, we describe the configuration limitations of the supported partitioning systems.

pSeries 690 Model 681

The following describes the configuration limitations for pSeries 690 Model 681:

- ▶ The pSeries 690 can be divided into as many as 32 logical partitions. System resources can be dedicated to each LPAR.
- ▶ Servers configured with Support Processor with Remote I/O Loop Attachment (FC 6404) can support up to 16 LPARs maximum.
- ▶ Servers configured with Support Processor with Remote I/O-2 (RIO-2) Loop attachment (FC 6418) can support up to 32 LPARs maximum.
- ▶ LPAR allocation, monitoring, and control is provided by the Hardware Management Console.
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of one processor and 1 GB of memory is required per LPAR.
- ▶ A minimum of 4 GB of system memory is recommended per LPAR.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR on an individual slot basis.
- ▶ Integrated Ultra3 SCSI controllers located in the 7040-61D I/O drawers may be individually allocated to an LPAR. These integrated SCSI adapters each support one 4-pack disk backplane.
- ▶ While it is not mandatory, consideration should be given to allocating one half of a 7040-61D I/O drawer for each LPAR. This would provide one 10-slot PCI or PCI-X planar, two integrated SCSI controllers, and two 4-Pack SCSI disk backplanes to the LPAR. This will help to ensure balanced I/O bandwidth for the LPAR configurations.

pSeries 670 Model 671

The following describes the configuration limitations for pSeries 670 Model 671:

- ▶ The pSeries 670 can be divided into as many as 16 logical partitions. System resources can be dedicated to each LPAR.
- ▶ p670 servers configured with Support Processor with Remote I/O Loop Attachment (FC 6404) can support up to 16 LPARs maximum.
- ▶ p670 servers configured with Support Processor with Remote I/O-2 (RIO-2) Loop Attachment (FC 6418) can support up to 16 LPARs maximum.
- ▶ LPAR allocation, monitoring, and control is provided by the Hardware Management Console.
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of one processor is required per LPAR.

- ▶ A minimum of 1 GB of system memory is required per LPAR. While it is not required, 4 GB of system memory is recommended per LPAR to ensure adequate memory is available.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR on an individual slot basis.
- ▶ Integrated Ultra3 SCSI controllers located in the 7040-61D I/O drawers may be individually allocated to an LPAR. These integrated SCSI adapters each support one 4-pack disk backplane.
- ▶ While it is not mandatory, consideration should be given to allocating one half of a 7040-61D I/O drawer for each LPAR. This would provide one 10-slot PCI or PCI-X planar, two integrated SCSI controllers, and two 4-Pack SCSI disk backplanes to the LPAR. This will help to ensure balanced I/O bandwidth for the LPAR configurations.

pSeries 655 Model 651

The following describes the configuration limitations for pSeries 655 Model 651:

- ▶ The pSeries 655 can be divided into four logical partitions (nonswitched or switched). System resources can be dedicated to each LPAR.
- ▶ LPAR allocation, monitoring, and control is provided by the HMC.
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of 4 GB of system memory is recommended per LPAR.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR on an individual slot basis.
- ▶ Integrated Ultra3 SCSI controllers may be individually allocated to an LPAR.

pSeries 650

The following describes the configuration limitations for pSeries 650:

- ▶ Static LPAR is supported with AIX 5L Version 5.1 or AIX 5L Version 5.2. Dynamic LPAR requires AIX 5L Version 5.2.
- ▶ A Hardware Management Console is required for all LPAR configurations.
- ▶ A fully configured pSeries 650 can support up to eight partitions. A partition must have at least a bootable SCSI or SSA or Fibre Channel adapter, and an Ethernet adapter/controller. Since all the PCI slots in p650 can be independently assigned to partitions, this means that four partitions with minimal I/O can be configured without requiring a separate 7311-D10 I/O drawer. A Model D10 I/O drawer can provide up to three additional minimal I/O partitions. To configure the maximum of eight partitions would require at least two Model D10 I/O drawers. The actual number of D10 I/O drawers required will depend on the customer's I/O requirements for each partition.

Note that both the integrated internal SCSI bus and the integrated external SCSI bus must be assigned to the same partition.

- ▶ All of the internal disk drives and media share the same SCSI bus and must be assigned to the same partition. That means that all but the first partition must be NIM installed and use network diagnostics, or can only be CD-ROM installed or use network diagnostics when the first partition is shut down. The user should not use the internal disk drives during install because drives could have data for the first partition and may be reassigned back to the first partition. It is recommended to use NIM and network diagnostics for all partitions except the first partition, or that the internal disk drives not be used for any partition.
- ▶ External disk drives must be provided for all but the first partition. AIX does not support diskless LPAR operation.
- ▶ When running LPAR with AIX 5L Version 5.1, fully populated processor cards cannot mix memory DIMMs with different capacities on the same card.

pSeries 630 Model 6C4/6E4

The following describe the configuration limitations for pSeries 630 Model 6C4 and 6E4:

- ▶ Feature number FC 9575 or FC 6576 is required on the Model 6C4 or 6E4 to run LPAR.
- ▶ AIX 5L Version 5.1 environments running in a p630 LPAR require a re-activation (reboot) of the affected partition when an LPAR is redefined.
- ▶ AIX 5L Version 5.2 environments running in a p630 LPAR can be redefined dynamically without a partition re-activation when non-critical resources and non-ISA resources are involved.
- ▶ The pSeries 630 Model 6E4 can be divided into three LPARs if the system has six I/O slots. Only two LPARs are possible in systems with four I/O slots. The Model 6C4 can be divided into as many as four LPARs. System resources can be dedicated to the LPARs.
- ▶ One integrated Ethernet controller and two I/O slots must be assigned to a single partition. Since the 7311-D20 I/O Expansion Drawer is not supported on the Model 6E4, two I/O slots are required for each additional partition, one for an Ethernet adapter and one for an external disk connection.
- ▶ LPAR allocation, monitoring, and control is provided by the Hardware Management Console (7315-C01).
- ▶ Each LPAR functions under its own instance of the operating system.
- ▶ A minimum of one processor is required per LPAR.

- ▶ One LAN adapter (minimum) per partition (HMC requirement). Two integrated Ethernet ports can be used to support two partitions.
- ▶ I/O adapters in PCI slots may be allocated to an LPAR in the following manner:
 - Slots 3 and 4 can be allocated to different partitions. To support four partitions, one 7311-D20 is required.
 - Slots 1 and 2 must be allocated to the same partition. This partition must also have:
 - Three integrated serial ports
 - Keyboard and mouse ports
 - Integrated dual Ultra3 SCSI ports
 - Any IDE media device
 - One of the integrated Ethernet ports
- ▶ If a SCSI DVD-RAM is installed in the system, the DVD-RAM can be defined to any partition for install purposes. If no DVD-RAM is in the system, a NIM connection must be available to all partitions but the first. For network diagnostics, a NIM connection is always required in all partitions but the first. A SCSI adapter is required to drive the DVD-RAM. The DVD-RAM cannot be installed in slot 1 or 2 if the DVD-RAM is to be shared.
- ▶ FC 2842 and 2843 are not supported in LPARs.
- ▶ Parallel port is not available in LPAR mode.
- ▶ At least one disk must be allocated to each LPAR.
- ▶ All DASD in the Model 6E4 and 6C4 must be allocated to one partition.
 - One 7311-D20 Expansion Drawer can be used to support two additional partitions on the Model 6C4. Two Model D20s support four partitions. Each partition attaching to a Model D20 requires a SCSI adapter in that partition if attaching to one of the two internal 6-packs.
 - The 7133 Serial Storage Devices can be used to support three additional partitions on the Model 6C4.
 - The 2104 Expandable Storage Devices can be used to support two additional partitions on the Model 6C4. Two 2104 devices support four partitions.
- ▶ A minimum of 768 MB of system memory is required per LPAR.
 - 768 MB is required for system overhead.
 - 1.5 GB will support two partitions.
 - 3 GB will support four partitions.

Note: While it is not required, 2 GB of system memory is recommended per LPAR to ensure adequate memory is available.

Logical partitioning support on pSeries 630 and pSeries 650

Due to the I/O structure on these pSeries models, some resources may have to be assigned (and moved) as a group.

On the pSeries 650 systems, the internal disks and the CD/DVD device that are on the same SCSI controller must be assigned together and moved together with dynamic logical partitioning (DLPAR) operations.

Note: Throughout this section, the term *CD/DVD devices* to refer to CD-ROM, DVD-RAM, and DVD-ROM devices.

On the pSeries 630 systems, PCI slot 1, PCI slot 2, internal Ethernet 2 (physical location code: U0.1-P1/E2), the internal SCSI disk subsystem (U0.1-P2/Z1), and ISA-based I/O (serial, keyboard, mouse) must be assigned to a single partition as a group and cannot be assigned separately to different partitions. These resources cannot be moved with DLPAR operations because they share a PCI host bridge with the Industry Standard Architecture (ISA) subsystem.

Note: The parallel ports on pSeries 630 are not supported in a partitioned environment.

16.1.5 Terminology used in partitioning

In this section, we explain the concepts and terminology used in logical partitioning on partitioning-capable pSeries servers.

Logical partitioned environment

The partitioning-capable pSeries servers support a logical partitioned environment that enables you to run multiple logical partitions concurrently. The maximum number of partitions that can concurrently run depends on the specific partitioning-capable pSeries server model. For example, both the pSeries 670 and pSeries 690 support up to 16 partitions running concurrently.

In a logical partition, an operating system instance runs with dedicated resources: processors, memory, and I/O slots. These resources are statically assigned to the logical partition. The total amount of assignable resources is limited by the physically installed resources in the system.

Because the first implementation of logical partitioning provided by AIX 5L Version 5.1 is static, you have to shut down every operating system instance in the targeted logical partitions to change between these two modes or to change the resource assignment of running logical partitions. Logical partitions that are not changed are unaffected.

With the release of AIX 5L Version 5.2, dynamic logical partitioning (also called DLPAR) extends the capabilities of static logical partitioning by allowing the dynamic reassignment of resources across partitions. Dynamic logical partitioning provides an improved solution by allowing users to dynamically move hardware resources between partitions without requiring a reboot of partitions.

All the partitioning-capable pSeries servers also support a special type of logical partition, called the Full System Partition. In a Full System Partition, only one operating system instance runs in the system. This instance has access to all the resources installed in the system. In Full System Partition, a partitioning-capable pSeries server acts as a conventional pSeries SMP server, except for the performance and enhanced reliability, availability, and serviceability (RAS) feature.

Partition isolation and security

From a functional point of view, applications are running inside partitions in the same way they run on a stand-alone pSeries machine. There are no issues when moving an application from a stand-alone server to a partition. The design of partitioning-capable pSeries servers is such that one partition is isolated from software running in the other partitions, including protection against natural software defects and even deliberate software attempts to break the partition barriers. It has the following security features:

- ▶ Protection against inter-partition data access

The design of partitioning-capable pSeries servers prevents any data access between partitions, other than using regular networks. This isolates the partitions against unauthorized access between partition boundaries.

- ▶ Unexpected partition crash

A software crash within a partition should not cause any disruption to the other partitions. Neither an application failure nor an operating system failure inside a partition interfere with the operation of other partitions.

- ▶ Denial of service across shared resources

The design of partitioning-capable pSeries servers prevents partitions from making extensive use of a shared resource so that other partitions using that resource become starved. This means that partitions sharing the same PCI bridge chips, for example, cannot occupy the bus indefinitely.

In this way, applications can be safely consolidated in partitions in a partitioning-capable pSeries server without compromising overall system security.

16.1.6 Four terms regarding memory

Because the word memory is overused in different contexts, we have to provide precise definitions of the four terms regarding memory (virtual, physical, real, and logical memory) before explaining the partitioning implementation details.

First, we introduce Figure 16-1 on page 698 to explain the relationship between virtual and physical memory.

The term *virtual memory* is used in many operating system environments to express the virtual memory function that enables the operating system to act as if the system were equipped with a larger memory size than it physically has. For example, there are two user processes, 1 and 2, in Figure 16-1 on page 698. Because each process should be isolated from the other processes, each has its own virtual memory address range, called the process address space. Each process address space is classified into several memory chunks called *segments* (shown as hatched rectangles in the figure). Each segment is again divided into small size memory chunks called *pages* (not shown in the figure). The page is the minimal allocation unit size of virtual memory. As shown in the Figure 16-1 on page 698, the process address space is partially filled, but is mostly vacant space.




Figure 16-1 Virtual and physical memory relationship

The term *physical memory* is also used in many operating system environments to express virtual memory function. Because all the virtual memory cannot sit in the physical memory in the system, only some portions of virtual memory are mapped to physical memory. The rest of the virtual memory is divided by page size, and each page can be mapped to a disk block in paging spaces, or still reside in a block of files in the file systems. This address translation is managed by the virtual memory manager (VMM) of the operating system using hardware components, such as the hardware page frame table (PFT) and translation look aside buffer (TLB).

The term *real memory* is often used to represent the physical memory, especially when discussing the VMM functionality in the kernel. The modifier *real* comes from the real addressing mode defined in some processor architectures (including PowerPC) where address translation is turned off. In a non-partitioned environment, because there is a one-to-one relationship between the real and physical memory, we can ignore the difference between these two terms in most cases.

The physical address space must encompass all addressable hardware components, such as memory cards, I/O ports, bus memory, and so on. Depending on the hardware implementation and restrictions, address ranges might need to be dispersed throughout the physical address space, which could result in a discontinuous physical memory address space. For example, if a PCI adapter device requires direct memory access (DMA), the device's DMA address is mapped on the specific physical memory address range by a PCI host bridge (PHB). Most VMMs of modern operating systems are designed to deal with the non-contiguous physical memory address.

However, operating systems require a certain amount of contiguous physical memory that can be addressed translate-off, typically for bootstrapping, in a non-partitioned environment.

In a partitioned environment, real and physical memories have to be distinctly distinguished, because the physical memory address, which previously meant the real memory address, is no longer being used in that way, because there is an extra level of addressing in a partitioned environment.

To support any operating system, including AIX and Linux, which require real mode code execution and the ability to present a real address space starting at 0 to each partition in the system, the *logical memory* concept is adopted. Logical memory is an abstract representation that provides a contiguous memory address to a partition (see Figure 16-2 on page 700). Multiple non-contiguous physical memory blocks are mapped to provide a contiguous logical memory address space. The logical address space provides the isolation and security of the partition operating system from direct access to physical memory, allowing the hypervisor to police valid logical address ranges assigned to the partition. The contiguous nature of the logical address space is geared more to simplifying the hypervisor's per-partition policing than it is due to an operating system requirement. The operating system's VMM handles the logical memory as if it was physical memory in a non-partitioned environment.




Figure 16-2 Non-contiguous mapping

We also use the following terms throughout this redbook:

Physical memory block (PMB)

The physically contiguous memory block unit size of 256 MB manipulated inside the global firmware. A PMB ID is unique throughout the system.

Logical memory block (LMB)

The memory block unit size of 256 MB seen from the partitions. An LMB ID is unique inside a partition. An LMB is associated to a PMB. Therefore, a partition can have the same LMB IDs as other partitions. However, the same LMB ID for different partitions are associated to different PMBs.

16.2 Partitioning implementation on pSeries servers

This section explains the partitioning implementation on pSeries servers in the the following sub-sections:

- ▶ Partitioning implementation
- ▶ Partitioning resources
- ▶ I/O device assignment considerations
- ▶ Service authority

16.2.1 Partitioning implementation

In a partitioned environment, each partition must have exclusive access to their assigned resources, but assigned resources also have to be strictly isolated from the other partitions. To implement these two demands, several system components have to work together to support the partitioning on partitioning-capable pSeries servers. We categorize these components into the three groups explained in the following sections:

- ▶ Hardware
- ▶ Firmware
- ▶ Operating system: AIX 5L Version 5.1
- ▶ Operating system: AIX 5L Version 5.2

Hardware

To support partitioning on partitioning-capable pSeries servers, in addition to the IBM Hardware Management Console for pSeries (HMC), the following hardware components have been newly developed or enhanced.

POWER4 processor

The POWER4 processor is a highly reliable and high-performance processor. Besides these characteristics, it also provides the following new functions to support a partitioned environment:

- ▶ Hypervisor call

The POWER4 processor supports a special form of instructions. These instructions are exclusively used by a new controlling firmware named *hypervisor*. If an operating system instance in a partition requires access to hardware, it first invokes the hypervisor using *hypervisor calls*. Hypervisor allows privileged access to an operating system instance for dedicated hardware facilities and includes protection for those facilities in the processor. We discuss the hypervisor functions in “Firmware” on page 703.

- ▶ Real mode offset (RMO) register

The POWER4 processor supports a real mode offset register (RMO). By utilizing the RMO, the processor enables a mapping between the logical memory of a partition and the physical memory. A logical memory address of zero in a partition is registered at a fixed offset in the physical memory address space. The RMO is the physical address that corresponds to the beginning of a partition's memory. This is the partition's logical address 0. This address offset is set in the RMO register when this partition is activated. The processor adds the value stored in the RMO to each logical address fetch and store made by code running in a partition.

The RMO address space only applies to Supervisor Real Mode Execution (MSR.DR/MSR.IR = 0). The RMO defines the low logical memory (starting at 0) that the operating system (Supervisor) can directly address in translation-off mode. Other logical address ranges (beyond the RMO) have no interaction or dependency on the real mode offset register, because they cannot be addressed in the translate-off mode.

- ▶ Real mode limit register (RML)

The POWER4 processor supports a real mode limit register (RML). By utilizing the RML, the processor enables you to limit the range of real mode addressing. The RML is the size of a partition's memory region that is accessed in real mode.

For further information about the POWER4 processor, refer to the following Redbooks:

- ▶ *IBM @server pSeries 670 and pSeries 690 System Handbook, SG24-7040*
- ▶ *The POWER4 Processor Introduction and Tuning Guide, SG24-7041*

Interrupt controller

In a Full System Partition, the interrupt controller that manages the peripheral interrupt requests to the processors works in a fashion similar to other pSeries SMP servers. In a partitioned environment, the interrupt controller supports multiple global interrupt queues, which can be individually programmed to send external interrupts only to the set of processors allocated to a specific partition. Therefore, the processors in a partition can only receive interrupt requests from devices inside their partition.

PCI host bridges

In a Full System Partition, the PCI host bridges (PHBs) control the PCI slots in the I/O drawers, as in conventional pSeries servers. The PHBs use translation control entry (TCE) tables for the I/O address to memory address translation in

order to perform direct memory access (DMA) transfers between memory and PCI adapters. The TCE tables are allocated in the physical memory.

In a partitioned environment, the hypervisor controls the DMA addressing to the partition memory for all I/O devices in all partitions. The hypervisor uses central TCE tables for all I/O devices, which are located outside of the memory of the partitions. The hypervisor can manage as many TCE tables as it needs to. For example, each PCI host bridge could have its own TCE table. The number of TCEs, and thus the number of TCE tables, needed is a function of the number of PCI host bridges and slots. The address mapping is protected on a per adapter basis. The PCI host bridges used in partitioning-capable pSeries servers support the control of the PCI adapter DMA by the hypervisor.

The key point is that a logical partition is only given a *window* of TCEs per I/O slot necessary to establish DMA mappings for the device in that slot. The hypervisor polices TCE access made by the operating systems to ensure they are to a window owned by the partition.

Error handling

Various system components have the ability to limit the impact of hardware errors on a single partition. Generally, this is achieved by turning most hardware error reporting into bad data packets that flow back to the requesting processor. In many cases, this will cause a machine check interrupt, which may or may not be recoverable within the partition. However, no other partitions are affected. The enhancement includes the enhanced error handling (EEH) on the PCI bus. For further explanation about EEH, refer to *IBM @server pSeries 670 and pSeries 690 System Handbook*, SG24-7040.

Service processor

All the partitioning-capable pSeries server models have an enhanced service processor (compared to existing pSeries models). The major enhancement of the service processor is a communication function with the IBM Hardware Management Console for pSeries (see Figure 16-5 on page 729).

Firmware

The support of partitioning on partitioning-capable pSeries servers requires a new firmware named hypervisor, partition Open Firmware, and Run-Time Abstraction Service (RTAS).

Hypervisor

The hypervisor firmware provides major additions to firmware functionality. It implements the following three major categories of service calls:

- ▶ Virtual memory management

Hypervisor becomes the only function that can update the address translation page tables in memory or the TCEs of the PHBs. In this way, hypervisor controls the physical memory locations that can be accessed from within a partition.

- ▶ Debug register and memory access

For the debug and dump environments, hypervisor provides controlled access to protected facilities and memory locations.

- ▶ Virtual TTY support

Hypervisor provides input/output streams for a virtual TTY device that can be used on the HMC.

The hypervisor is a passive object loaded into the first PMB in a partitioned environment. It is loaded only when the system is running in a partitioned environment and does not reserve a processor resource for itself. Hypervisor only runs when a partition needs a service executed on its behalf, such as creating a page table entry. Hypervisor can be thought of as a call-back library used as any partition requires. Care has been taken to minimize the number of instructions required to implement the call-backs, so in most cases, AIX performance is identical for AIX in a non-partitioned environment where call-backs are not made, versus AIX in a partitioned environment where call-backs are required.

The hypervisor resides outside of the partition system memory in the first PMB at the physical address zero. This first PMB is not usable by any of the partition operating systems in a partitioned environment.

Open Firmware

A partitioning-capable pSeries server has one instance of Open Firmware both in the partitioned environment and the Full System Partition. Open Firmware has access to all devices and data in the system. Open Firmware is started when the system goes through a power-on reset. Open Firmware, which runs in addition to the hypervisor in a partitioned environment, runs in two modes: global and partition. Global and partition Open Firmware share the same firmware binary stored in the flash memory.

In a partitioned environment, the partition Open Firmware runs on top of the global Open Firmware instance. The partition Open Firmware is started when a partition is activated. Each partition has its own instance of the partition Open Firmware, and it has access to all the devices assigned to that partition, but has no access to devices outside the partition in which it runs. Partition firmware resides within the partition memory, but is replaced when AIX takes control; it is just needed for the time necessary to load AIX into the partition system memory.

The global firmware resides with the hypervisor firmware in the first 256 MB of the physical memory.

The global Open Firmware includes the partition manager component. The partition manager is an application in the global Open Firmware that establishes partitions and their corresponding resources, such as CPU, memory, and I/O slots, which are defined in partition profiles. The partition manager manages the operational partitioning transactions. The partition manager responds to commands from the service processor external command interface that originate in the application running on the HMC.

The partition profiles are stored in and retrieved from nonvolatile random access memory (NVRAM) by system firmware upon requests sent from the HMC. After the profiles are set up, the system will automatically return to this configured state on a power-on, even if the HMC is unavailable. The NVRAM also provides separate address spaces, called *slots*, to store up to several partitions. These slots, or partition IDs, are numbered from 1 to the supported maximum partition number.

One non-virtual, hardware password (PAP) controls access to systems management services (SMS), per system) which is generally presented on the virtual terminal window. This is consistent with the system administrator having access to the HMC.

To confirm the current firmware level, you can use the `lscfg` command as follows:

```
# lscfg -vp | grep -p 'Platform Firmware:'
Platform Firmware:
  ROM Level.(alterable).....RH030307
  Version.....RS6K
  System Info Specific.(YL)...U1.18-P1-H2/Y1
  Physical Location: U1.18-P1-H2/Y1
```

This example shows firmware level RH030307.

To download and recognize the firmware level, you can use the following link.

<http://techsupport.services.ibm.com/server/mdownload/>

Run-Time Abstraction Services (RTAS)

RTAS presents the same platform service calls (with a few exceptions) that are presented in a non-partitioned environment, but have some underlying implementation changes to properly handle multiple AIX images, including:

- ▶ RTAS calls are only serialized within a partition.

In general, RTAS operations are restricted to only those resources dedicated to that partition, with an error code return for invalid requests.

- ▶ Multiple virtual operator panels for all partitions.

The information provided by operator panels in a traditional pSeries server are represented on the HMC on a per-partition basis.

- ▶ Per-partition time-of-day clock values.

Time-of-day (TOD) is virtualized for each partition (including Full System Partition). Each partition can set its own time and date using the AIX **date mdddhMMYYyy** command.

- ▶ Restricted access to the per-partition NVRAM areas.

Each partition has its own segment of NVRAM for the storage of its configuration variables, including a unique boot list for every partition. There is also a unique segment of NVRAM for when the system is in Full System Partition, with its own boot list. For example, there are up to 16 partition boot lists and a seventeenth Full System Partition boot list on the pSeries 670 and pSeries 690 equipped with 7040-61D I/O drawers.

- ▶ Partition reset capabilities.

Previous pSeries servers had a service processor-based serial port snoop function that allowed remote reset of an unresponsive AIX image. The service processor would snoop the serial port data stream (which is the AIX console), but when AIX is not longer reachable through the keystrokes, and when a certain special command sequence was seen, the service processor would reset the system.

On partitioning-capable pSeries servers, with the HMC, each partition has its own, very powerful, reset capabilities: a soft reset that causes the partition operating system to get a PowerPC reset interrupt, and a hard reset that is the equivalent of a virtual power off of a partition. For the hard reset, no matter how disabled the partition operating system is, the hard reset will bring all the processors out of the partition and back to the global firmware partition manager so that the partition is ready to be reactivated.

To confirm the current Open Firmware RTAS version, you can use the **lsattr** command as follows:

```
# lsattr -El sys0 | grep rtasversion
rtasversion 1                Open Firmware RTAS version                False
```

This example shows RTAS Version 1.

Operating system: AIX 5L Version 5.1

The operating system in a partition must use the hypervisor calls in place of direct access to the hardware and address mapping facilities in conventional pSeries machines. An operating system without partition-enabled functions does not work in a partition.

If you are going to use AIX as a partition-enabled operating system, you will need to use AIX 5L Version 5.1 with the 5100-01 Recommended Maintenance Level plus APAR IY22854. From the operational point of view, there are a few noticeable differences in AIX when running inside a partition, as shown in Figure 16-3 on page 709:

- ▶ There is no physical console on the partition, unless you assign it explicitly. The built-in native serial ports¹ on the partitioning-capable pSeries server can be assigned together only to one partition at the same time. To provide an output for console messages, and also for diagnostic purposes, the firmware implements a virtual TTY, called virtual terminal, that is seen by AIX as a standard TTY device. Its output is streamed to the HMC. The AIX diagnostics subsystem uses this virtual TTY as a system console.
- ▶ Certain platform operations are constrained in partitions. For example, in a non-partitioned environment, platform firmware updates can be performed from AIX by the root user. Because firmware updates may affect all partitions in a partitioned environment, the administrator has the ability to specify that a particular partition has this authority. Within that partition, firmware updates will work in the same way as they do for non-partitioned environment. See 16.2.4, “Service authority” on page 724 for more information.

Besides these considerations, AIX runs inside a partition the same way it runs on a stand-alone pSeries server. No difference is observed either from the application or the administrator’s point of view.

Note: The 32-bit AIX kernel supports up to 96 GB of physical memory size and 32-way processors, regardless of whether it is in a partition or in a Full System Partition.

Partitions are very transparent to AIX applications. In fact, third-party applications only need to be certified for a level of AIX that runs in a partition, and not for the partitioned environment itself. In this way, a partition on a partitioning-capable pSeries server can be viewed as just another pSeries hardware platform environment.

¹ The pSeries 655 Model 651 has no built-in native serial port.




Figure 16-3 Interaction of AIX and firmware in a partition

Paging performance in a partitioned environment

An operating system in a partition has slightly less page table management performance, because it must use hypervisor services for page table management. An operating system in a Full System Partition has full use of all the system memory and *native* virtual memory management performance. In high volume paging environments, system performance is slightly less than native performance. In normal paging environments, there is no observable difference in performance.

Fast reboot in a partitioned environment

Rebooting an operating system instance in a partition is much faster than a full system reboot of a comparable conventional pSeries system because less hardware initialization is required.

Partition reboots are merely a re-establishment of the pSeries Open Firmware OS boot loader environment and, by nature, are very quick. A Full System Partition reboot repeats all the hardware initialization phases of the processors, caches, and memory. These phases are done by the service processor, and the I/O drawers and I/O adapters are done by the system firmware. When configuring all system resources in a single partition, hypervisor remains resident in memory. This enables the extremely rapid re-establishment of the boot environment, but requires the reservation of the first PMB by the hypervisor (see “Reserved memory regions in a partitioned environment” on page 715).

Full System Partition reboots still have full system initialization phases and are almost comparable to traditional pSeries reboot times.

Affinity logical partitions

On the pSeries 670 and pSeries 690, another type of partition, called affinity logical partitions, have been supported with the following configuration:

- ▶ AIX 5L Version 5.1 with 5100-02 Recommended Maintenance Level or later
- ▶ HMC software Release 2 Version 1 or later
- ▶ April 2002 system microcode update or later

Affinity logical partitions are useful in order to run CPU and memory-intensive applications, typically found in the High Performance Computing (HPC) area, on these server models

Large page support

Although AIX uses a fixed size 4 KB page regardless of the release, an enhancement called *large page* was introduced in AIX 5L POWER for Version 5.1 with 5100-02 Recommended Maintenance Level. Exploiting the large page function, memory intensive applications, such as HPC and relational database management system (RDBMS) applications, can get performance benefit, because less overhead is required by the hardware and VMM. For further information about large pages, please refer to the *AIX Support for Large Pages* white paper, found at:

http://www.ibm.com/servers/aix/whitepapers/large_page.html

Operating system: AIX 5L Version 5.2

In addition to all the enhancements and components provided in the former releases, AIX 5L Version 5.2 provides many enhancements, which are explained in the *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

Dynamic logical partitioning

Starting with Version 5.2, AIX supports dynamic logical partitioning. Dynamic logical partitioning is a function to allow a partition whose resources are dynamically added and removed without requiring a partition reboot. Both 32- and 64-bit kernels running in a partition (except for the Full System Partition) support the dynamic logical partitioning function.

The support of dynamic logical partitioning also provides the following features:

- ▶ Dynamic processor deallocation with the minimum of two processors
AIX supports the dynamic processor deallocation² function starting with Version 4.3.3, which can dynamically take a processor offline when an internal threshold of recoverable errors on the processor is exceeded. Before AIX 5L Version 5.2, an SMP system or a partition must have at least three processors, while an SMP system or a partition installed with AIX 5L Version 5.2 requires only two processors, in order to support dynamic processor deallocation.
- ▶ Capacity Upgrade on Demand (CUoD)
CUoD is a feature that allows customer to dynamically activate pre-installed but inactive processors by purchasing the license key without requiring a system reboot.
For further information about CUoD, please refer to *IBM @server pSeries 670 and pSeries 690 System Handbook*, SG24-7040.

Partition on Demand

The 10/2002 system microcode update provides the ability to plug adapters into empty slots, define a new partition at the HMC, and activate and boot a previously nonexistent partition, while other partitions are running, without the power recycle of the system. This functionality, called *Partition on Demand*, obviously needs empty PCI slots.

Previously, you could not do this because the partition boot did not change the power state of PCI slots. However, you could power on PCI slots on existing partitions with AIX PCI Hot Plug, but not *spontaneously* create a new partition, including a new boot device, and activate that new partition with all the other partitions running. One limitation is that when you define your partition, you will not see the type of PCI cards (SCSI, network, and so on) present on the HMC, because the PCI slots are still powered off. You can also do a DLPAR I/O slot remove and then build a new partition.

You can always power off the central electronics complex (CEC) or halt partitions, change I/O assignments, and restart, but this is quite severe on the running systems and not as common. With DLPAR in the 10/2002 system microcode update and AIX 5L Version 5.2, you can drop processors, memory, and I/O from an existing partition or set of partitions using DLPAR remove, and then spawn one or more new partitions that did not exist before. Spontaneous partitions up to a memory size of 16 GB are supported. With a memory size greater than 16 GB, two contiguous LMBs are required for the spontaneous partition's page table, and all the LMBs shed by the partition or partitions could all

² The dynamic processor deallocation function is also known as CPU Guard.

be discontinuous. Although it is possible to create a spontaneous partition with a memory size greater than 16 GB, Partition on Demand does not guarantee that it is always possible.

16.2.2 Partition resources

The logical partitioning function on partitioning-capable pSeries servers allows you to assign processors, physical memory, and I/O devices to partitions. In the following section, we explain the rules of resource assignment.

Partition and system profiles

The information about resources assigned to a partition is stored in a partition *profile*. Each partition can have multiple partition profiles. By switching from one partition profile to another, you can change how resources are assigned. For example, you can assign relatively small resources for small online transactions on weekdays, and large resources for high-volume batch transactions on weekends. In a static environment, to switch partition profiles, you have to shut down the operating system instance that is running in the partition and stop the partition (deactivate). You can also define a system profile (for administrative purposes) as an optional task. By using a system profile, you can power on multiple partitions in a specific order in one operation.

A description of the two profiles follows:

Partition profile A partition profile stores the information about the assigned resources for a specific partition, such as processor, memory, and I/O devices. Each partition must have a unique name and at least one partition profile. A partition can have several partition profiles, but it reads only one partition profile when it is started (*activated*). You select a partition profile when you activate the partition; otherwise, the default partition profile³ is used. You can designate any partition profile as the default partition profile.

System profile A system profile provides a collection of partition profiles that should be started at the same time. The partition profiles are activated in the order of the list defined in the system profile. You can also use system profiles to start the Full System Partition.

Both types of profiles are stored in the NVRAM of the partitioning-capable pSeries server. Although you can create many partition profiles and system partition profiles, because both types of profiles share the same memory area in

³ If you have only one partition profile for a partition, it is always the default partition profile.

the NVRAM, the actual number you can create depends on your profile configuration.

We summarize the relationship among partitions, partition profiles, and system profiles in Figure 16-4. In this figure, partitions A, B, and C have three, one, and two partition profiles, respectively. Each partition has the default partition profile represented with a check mark. The system profile X is associated with partition profiles A1, B1, and C2; also, the system profile Y is associated with partition profiles A1 and C1. Keep in mind the following points:

- ▶ You do not have to associate all the partition profiles with system profiles. In this example, the partition profiles A2 and A3 are not associated with any system profile.
- ▶ It is possible to associate a partition profile to multiple system profiles. In this example, the partition profile A1 is associated with system profile X and Y.

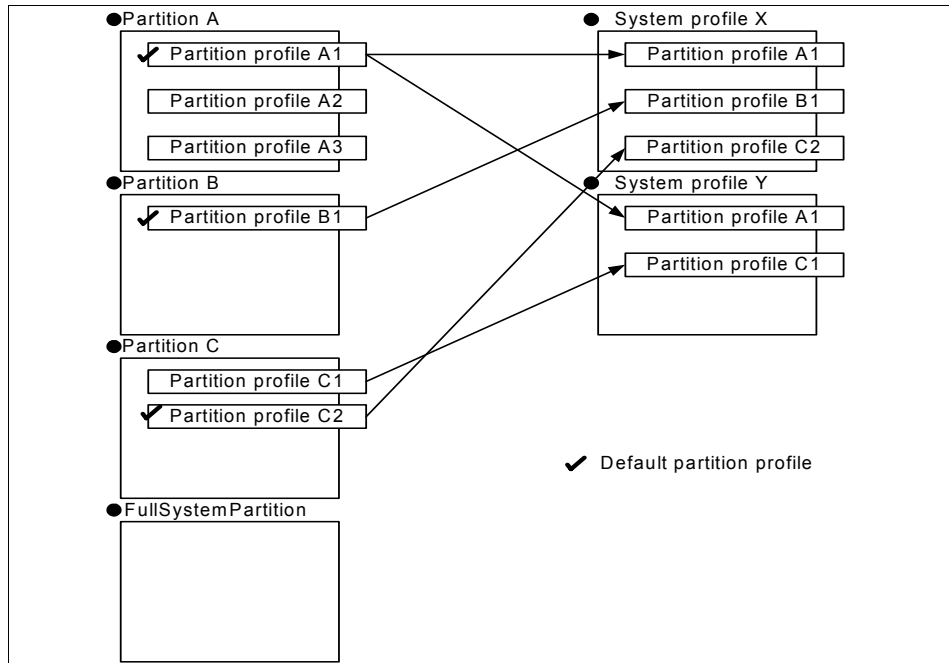


Figure 16-4 Partitions, partition profiles, and system profiles

To create partition profiles and system profiles, you have to use the IBM Hardware Management Console (HMC) for pSeries.

Three assignable resource types

A partition profile stores the information of the three types of resources: CPU, memory, and I/O slots. Partition profiles store the information of the specific PCI slots assigned in the I/O drawers where the I/O devices are possibly plugged in.

Note: Partition profiles do not store the information about specific I/O devices and PCI adapters.

Processors

Each installed and configured processor in the partitioning-capable pSeries server can be assigned to a partition. You do not have to specify the precise location of the assigned processors in the partition profile, because the system selects the resources automatically. At least one processor must be assigned to each partition. Sharing processors between multiple active partitions is not possible.

Memory

In a partitioned environment, some of the physical memory areas are reserved by several system functions to enable partitioning in the partitioning-capable pSeries server (see “Reserved memory regions in a partitioned environment” on page 715). You can assign unused physical memory to a partition. You do not have to specify the precise address of the assigned physical memory in the partition profile, because the system selects the resources automatically.

The minimum amount of physical memory for each partition is 256 MB.⁴ You can assign further physical memory to partitions in increments of 256 MB.

The AIX Virtual Memory Manager (VMM) manages the logical memory within a partition as it does the real memory in a stand-alone pSeries server. The hypervisor and the POWER4 processor manage access to the physical memory.

I/O slots

I/O devices are assignable to partitions on a PCI slot (physical PCI connector) basis. This means that it is not the PCI adapters in the PCI slots that are assigned as partition resources, but the PCI slots in which the PCI adapters are plugged.

To install an operating system, you have to assign at least one device adapter, typically a SCSI adapter, that is able to boot the operating system, and an adapter to access the install media (see “Boot devices” on page 719).

⁴ On a pSeries 670 and pSeries 690 with a firmware level earlier than the 10/2002 system microcode update, the minimum amount of physical memory for each partition is 1 GB.

Once installed, you need at least one device adapter⁵ connected to the boot disk or disks. For application use and system management purposes, you also have to assign at least one network adapter.

You can allocate slots in any I/O drawer on the system. We recommend that you assign more PCI slots than required for the number of adapters in the partition, even if these PCI slots are not populated with PCI adapters. This provides you with the flexibility to add PCI adapters into the empty slots of an active partition, using the PCI Hot Plug insertion/removal capability.

For detailed information about the I/O slots assignment, see 16.2.3, “I/O device assignment considerations” on page 717.

Three kinds of values for resource assignment

In a partition profile, you have to specify three kinds of values for each resource. For CPU and memory, you have to specify *minimum*, *desired*, and *maximum* values. For I/O slots, you have to specify the *required* and *desired* values.

If any of the three types of resources cannot satisfy the specified minimum and required values, the activation of a partition will fail. If the available resources satisfy all the minimum and required values, but do not satisfy desired values, the activated partition will get as many of the resources as are available.

The maximum value⁶ is used to limit the maximum CPU and memory resources when dynamic logical partitioning operations are performed on the partition.

Note: If you are going to install operating systems that do not support dynamic logical partitioning, you should specify the same values for both the desired and maximum values.

Reserved memory regions in a partitioned environment

In a partitioned environment, some of the physical memory regions are reserved by several system functions to enable partitioning on partitioning-capable pSeries servers. Before understanding the mapping between the logical memory address of a partition and the physical memory address, you have to consider the following memory regions:

- ▶ Hypervisor
- ▶ Partition page tables
- ▶ Translation control entry (TCE) tables

⁵ AIX installed in a partition can boot from SCSI, SSA, and Fibre Channel attached disks.

⁶ The maximum value is not shown and is unavailable if the HMC software level is earlier than Release 3 Version 1.

These three memory regions are not usable for the physical memory allocation of the partition.

Summary

The maximum possible number of partitions less or greater than 16 GB with the total memory size are summarized in Table 16-3. This table also provides the reserved and allocatable memory to partitions.

Table 16-3 Physical memory size and number of allocatable partitions

Total memory	Approximate memory overhead	Approximate usable partition memory	Maximum number of partitions <ul style="list-style-type: none"> ▶ AIX 5L Versions 5.1 and 5.2 ▶ Pre-10/2002 firmware (Notes 1, 2, and 3) 	Maximum number of partitions <ul style="list-style-type: none"> ▶ AIX 5L Version 5.1 ▶ Post-10/2002 firmware (Notes 1, 2, and 4) 	Maximum number of partitions <ul style="list-style-type: none"> ▶ AIX5L Version 5.2 ▶ Post-10/2002 firmware (Notes 2 and 5)
4 GB	.75 - 1 GB	3 - 3.25 GB	2 and 0	13 and 0	13
8 GB	.75 - 1 GB	7 - 7.25 GB	6 and 0	16 and 0	16
16 GB	.75 - 1 GB	15 - 15.25 GB	14 and 0	16 and 0	16
24 GB	1 - 1.25 GB	22.75 - 23 GB	16 and 0	16 and 0	16
32 GB	1 - 1.25 GB	30.75 - 31 GB	16 and 0	16 and 0	16
40 GB	1.25 - 1.75	46.25 - 46.75 GB	16 and 1	16 and 1	16
48 GB	1.25 - 1.75 GB	46.25 - 46.75 GB	16 and 1	16 and 1	16
64 GB	1.5 - 2 GB	62 - 62.5 GB	16 and 2	16 and 2	16
80 GB	2G - 2.5 GB	93.5 - 94 GB	16 and 3	16 and 3	16
96 GB	2 - 2.5 GB	93.5 - 94 GB	16 and 4	16 and 4	16
128 GB	2.5 - 3.5 GB	124.5 - 125.5 GB	16 and 6	16 and 6	16
160 GB	3.5 - 4.5 GB	187.5 to 188.5 GB	16 and 8	16 and 8	16
192 GB	3.5 - 4.5 GB	187.5 to 188.5 GB	16 and 10	16 and 10	16
224 GB	5 - 6 GB	250 to 251 GB	16 and 12	16 and 12	16
256 GB	5 - 6 GB	250 to 251 GB	16 and 14	16 and 14	16

Notes:

1. In column 4 and 5, a difference is made between partitions with memory less or equal to 16 GB and greater than 16 GB, respectively.
2. All partition maximums are subject to availability of sufficient processor, memory, and I/O resources to support that number of partitions. For example, a system with only eight processors can only support a maximum of eight partitions.
3. These rules apply to systems running partitions with any version of AIX or Linux, if the firmware and HMC release levels are earlier than the 10/2002 release level.
4. These rules apply to systems running partitions with AIX 5L Version 5.1, if the firmware and HMC release levels are at the 10/2002 release level or later. The HMC partition profile option for Small Real Mode Address Region should not be selected for AIX 5L Version 5.1 partitions. These numbers reflect the maximum when running only AIX 5L Version 5.1 partitions, but AIX 5L Version 5.1 and AIX 5L Version 5.2 partitions can be mixed, and may allow for additional partitions to be run (up to the maximum of 16).
5. These rules apply to systems running partitions with AIX 5L Version 5.2 (or later) or Linux, if the firmware and HMC release levels are at the 10/2002 release level or later. The HMC partition profile option for the Small Real Mode Address Region check box should be selected for these partitions.

The following points summarize the physical memory allocation:

- ▶ For systems with 16 GB or less of physical memory installed, this rule is valid:
The maximum number of partitions = Total memory (in GB) - 2.
- ▶ You have to install more than 32 GB physical memory on the partitioning-capable pSeries server to activate AIX 5L Version 5.1 partitions greater than 16 GB. Because the 16 GB real mode region should be aligned on physical address 16 GB, 32 GB, 48 GB, and so on, in 32 GB memory configuration, the physical memory addresses 16 GB to 32 GB is partially used by TCE and cannot be used to allocate the 16 GB real mode region.

16.2.3 I/O device assignment considerations

Assignment of I/O slots to partitions is a relatively easy task once you understand the following considerations. The hardware architecture of each partitioning-capable pSeries server model also influences some aspects of it.

For detailed information about each of the partitioning-capable pSeries server models, please refer to the following publications:

- ▶ For the pSeries 630 Model 6C4 and pSeries 630 Model 6E4:
pSeries 630 Models 6C4 and 6E4 Technical Overview and Introduction, found at:
<http://techsupport.services.ibm.com/server/library>
- ▶ For the pSeries 650 Model 6M2:
pSeries 650 Model 6M2 Technical Overview and Introduction, found at:
<http://techsupport.services.ibm.com/server/library>
- ▶ For the pSeries 670 and pSeries 690:
IBM @server pSeries 670 and pSeries 690 System Handbook, SG24-7040

Media devices

If your installation media is removable media (CD-ROM, DVD-RAM, 4 mm tape, and so on), the corresponding devices should be configured. However, the way of configuring removable media devices depends on the hardware architecture of partitioning-capable pSeries server, as described here:

- ▶ pSeries 630 Model 6C4 and pSeries 630 Model 6E4
An internal IDE CD-ROM (FC 2633) or DVD-ROM (FC 6634) device can be configured on these models. However, there are several device-reassignment operations required on the HMC in order to use these devices as the installation media device on these models if you run multiple partitions. You can also configure the following SCSI-attached removable media devices on these models⁷:
 - FC 2623: DVD-RAM drive (4.7 GB per surface)
 - FC 6120: 8 mm 80/160 GB tape drive
 - FC 6134: 8 mm 60/150 GB tape drive
 - FC 6158: 4 mm 20/40 GB tape drive
- ▶ pSeries 650
The pSeries 650 Model 6M2 supports the following SCSI-attached removable media devices:
 - FC 2635: 16/48X DVD-ROM auto-docking module
 - FC 2629: 4.7 GB R/W DVD-RAM auto-docking module
 - FC 2628: 40X CD-ROM auto-docking module
 - FC 6169: 8 mm 80/160 GB auto-docking module

⁷ These features require a single SCSI adapter (FC 6203 with 4260 2-drop connector cable) inserted in a PCI slot within the machine.

- FC 6131: 8 mm 60/150 GB auto-docking module
- FC 6185: 4 mm 20/40 GB auto-docking module

However, there are several device-reassignment operations required on the HMC in order to use these devices as the installation media device on the pSeries 650 Model 6M2. Because the SCSI controller connected to devices in the auto-docking media bays is shared with disk drives in the internal four hot-swappable disk drive bays, if a partition is using these disk drives as its boot device, you must shut down and power off that partition before reassigning the SCSI controller to another partition.

▶ pSeries 655 Model 651

The pSeries 655 Model 651 does not support any removable media devices. You can install AIX 5L Version 5.1 using PSSP⁸ (requires control workstation) on this model.

▶ pSeries 670 and pSeries 690

Configure devices in the media drawer and assign the SCSI adapter connected to the media drawer to the partition.

If your installation media is on the network, one of the following network adapters must be assigned to the partition:

- ▶ Ethernet
- ▶ Token ring

Boot devices

Each partition requires its own separate boot device. Therefore, you must assign at least one boot device and a corresponding adapter per partition. The partitioning-capable pSeries server models support boot devices connected with SCSI, SSA, and Fibre Channel adapters.

The following describes boot device considerations:

▶ pSeries 630 Model 6C4

The pSeries 630 Model 6C4 can have up to four internal SCSI disk drives in the 4-pack disk bay. However, these disks are connected to one internal SCSI controller, so they only can be assigned to a partition. The other partitions must be assigned the boot adapter and disk drive from the following options:

- A boot adapter inserted in one of four PCI-X slots in the system. A bootable external disk subsystem is connected to this adapter.
- A portable SCSI adapter (which can have various features) is inserted in the PCI-X slot 7 in a 7311-D20 I/O drawer connected to the system. The

⁸ Parallel System Support Program

adapter is connected to one of 6-pack disk bays of drawer that houses the boot disk drive.

- A boot adapter inserted in one of seven PCI-X slots in a 7311-D20 I/O drawer connected to the system. A bootable external disk subsystem is connected to this adapter.

Note: The pSeries 630 Model 6C4 supports up to two 7311-D20 I/O drawers.

▶ pSeries 630 Model 6E4

The pSeries 630 Model 6C4 can have up to four internal SCSI disk drives in the 4-pack disk bay. However, these disks are connected to one internal SCSI controller, so they only can be assigned to a partition. The other partitions must be assigned the boot adapter and disk drive by the following:

- A boot adapter inserted in one of four PCI-X slots in the system. A bootable external disk subsystem is connected to this adapter.

Note: The pSeries 630 Model 6E4 does *not* support any I/O drawers.

▶ pSeries 650

The pSeries 650 Model 6M2 can have up to four internal SCSI disk drives in the 4-pack disk bay. However, these disks are connected to one internal SCSI controller, so they only can be assigned to a partition. The other partitions must be assigned the boot adapter and disk drive from the following options:

- A boot adapter is inserted in one of seven PCI-X slots in the system. A bootable external disk subsystem is connected to this adapter.
- A boot adapter is inserted in one of six PCI-X slots in a 7311-D10 I/O drawer connected to the system. A bootable external disk subsystem is connected to this adapter.

Note: The pSeries 650 Model 6M2 supports up to eight 7311-D10 I/O drawers. The 7311-D10 I/O drawer does *not* have any internal disk drives. Therefore, bootable external disk subsystems are mandatory in order to configure and run multiple partitions on the pSeries 650 Model 6M2.

► pSeries 655 Model 651

The pSeries 655 Model 651 can have up to two hot-swappable internal SCSI disk drives. However, these disks are connected to one internal SCSI controller, so they only can be assigned to a partition. Another partition must be assigned the boot adapter and disk drive from the following options:

- A boot adapter inserted in one of three PCI slots in the system. A bootable external disk subsystem is connected to this adapter.
- An internal disk drive inserted in one of the 4-pack disk bays on a 7040-61D I/O drawer and the SCSI controller on the drawer.

The 7040-61D I/O drawer can have up to 16 internal SCSI disk drives in the four 4-pack disk bays. Each of the disk bays is connected to a separate internal SCSI controller on the drawer.

- A boot adapter is inserted in one of 20 PCI slots in a 7040-D61 I/O drawer connected to the system. A bootable external disk subsystem is connected to this adapter.

Note: The pSeries 655 Model 651 supports one 7040-61D I/O drawer.

► pSeries 670 and pSeries 690

Partitions must be assigned the boot adapter and disk drive from the following options:

- An internal disk drive is inserted in one of the 4-pack disk bays on I/O drawer and the SCSI controller on the drawer.

The 7040-61D I/O drawer can have up to 16 internal SCSI disk drives in the four 4-pack disk bays. Each of the disk bays is connected to a separate internal SCSI controller on the drawer.

- A boot adapter is inserted in one of 20 PCI slots in a 7040-61D I/O drawer connected to the system. A bootable external disk subsystem is connected to this adapter.

Note: The pSeries 670 supports up to three 7040-61D I/O drawers, and the pSeries 690 supports up to eight. The minimum hardware configurations of these models require at least one I/O drawer.

You should select the adapter of the boot device from the PCI slot of the system or the first I/O drawer (on pSeries 670 and pSeries 690) if the system is running in a Full System Partition, because the system can quickly find the boot device. In a partitioned environment, the placement of the boot adapter does not affect the speed of the partitions' boot.

Network devices

It is mandatory to assign a network adapter to each partition. In addition to providing network access to client systems of a partition, the connection is also needed to provide the capability to manage the operating system and the applications in the partition remotely, either with a Telnet session or a graphical user interface, such as the Web-based System Manager. An Ethernet network connection between partitions and the HMC must be available if you want to use one of the following services:

- ▶ Service Agent
- ▶ Service Focal Point
- ▶ Inventory Scout
- ▶ Dynamic logical partitioning

These services communicate over the TCP/IP network between the partitions and the HMC.

Native Industry Standard Architecture (ISA) devices

All pSeries server models are equipped with several natively equipped ISA devices, such as native serial ports, a diskette drive, and keyboard and mouse ports, in order to support minimum hardware startup and diagnostics requirement. These ISA devices have the following characteristics when assigning resources to a partition:

- ▶ ISA devices are typically equipped on the system planner of the system (except for the pSeries 670 and pSeries 690) and they are connected by a single PCI bus through the PCI-ISA bridge chip. Therefore, these ISA devices can be assigned to only one partition as a group.
- ▶ ISA devices are not supported by dynamic logical partitioning.

The following list shows ISA devices on the partitioning-capable pSeries servers:

- ▶ pSeries 630 Model 6C4 and pSeries 630 Model 6C4

These models have the following ISA devices:

- A diskette drive (if configured)
- An IDE CD-ROM or DVD-ROM drive (if configured)
- Three native serial ports (S1R, S2, and S3)
- Keyboard and mouse ports

Note: The parallel ports on these models are not supported in a partitioned environment.

▶ pSeries 650

The pSeries 650 Model 6M2 has the following ISA devices:

- A diskette drive
- Four native serial ports (S1, S2, S3, and S4)
- Keyboard and mouse ports

▶ pSeries 655 Model 651

The pSeries 655 Model 651 has no assignable ISA devices.

▶ pSeries 670 and pSeries 690

These models have the following ISA devices:

- Two native serial ports (S1 and S2) in the primary I/O book
- A diskette drive in the media drawer

These devices are shown as Group_XXX in the I/O selection when creating a partition profile and the corresponding explanation.

Console devices

The HMC provides one virtual TTY console, called virtual terminal, for each partition, which removes most of the need for partition access to native serial ports. However, this virtual terminal was designed for limited purposes, such as installation of AIX and running diagnostics. If you need direct console access without using the network, there are two possibilities:

▶ Serial console

To connect an ASCII terminal or a terminal server to a partition as a serial console, you have to assign an 8-port (FC 2943) or 128-port (FC 2944) serial adapter to the partition.

▶ Graphics console

A graphics console is available on a partition by configuring the following features on the partition:

- A graphics adapter (FC 2848) with a graphics display
- A USB keyboard and mouse adapter (FC 2737) with a USB keyboard and a USB mouse attached.

Only one graphics console is supported per partition.⁹ The graphics console is functional only when AIX is running. For any installation or service processor support functions, you have to use the virtual terminal function on the HMC.

⁹ Up to eight partitions can have graphics consoles on the pSeries 670 and pSeries 690.

High availability

You should place redundant devices of a partition in separate I/O drawers for highest availability. For example, if two Fibre Channel adapters support multipath I/O to one logical unit number (LUN), and if one path fails, another path using another adapter in another I/O drawer is automatically chosen by the device driver.

Some PCI adapters do not have enhanced error handling (EEH) capabilities built into their device drivers. If these devices fail, the PCI host bridge (PHB) they are placed in and the other adapters in this PHB will be affected. Therefore, it is strongly recommended that you place all adapters without EEH capabilities on their own PHB, and not assign non-EEH adapters on the same PHB to different partitions.

For detailed information about EEH and supported PCI adapters on the pSeries servers, please refer to *PCI Adapter Placement References*, SA38-0538.

16.2.4 Service authority

You can give one of the partitions in a partitioning-capable pSeries server the *service authority* attribute. Service authority enables this partition, if running the AIX operating system, to perform system firmware updates or to set system policy parameters. Firmware updates also can be done from the service processor menus. Firmware updates are done at the system level, not on a per-partition basis.

A partition with service authority can perform firmware updates without having to power off the managed system. All other partitions must be shut down before the firmware update is initiated. The partition that has service authority must also have access to the firmware update image. If the firmware update image is provided on diskette, the diskette drive must belong to the partition that has service authority. If you are downloading the firmware update from the network, download it to the partition with service authority.

In the Full System Partition, you do not have to take additional steps to prepare for firmware updates.

16.3 Dynamic logical partitioning

Starting from AIX 5L Version 5.2, AIX supports dynamic logical partitioning (DLPAR). DLPAR is a function to allow a partition whose resources are dynamically added and removed without requiring a partition reboot.

The function enables the partitioning-capable pSeries server models to be used in the strategic autonomic computing infrastructure by dynamically shifting resources among partitions on a single system.

16.3.1 Dynamic logical partitioning overview

DLPAR supports the following dynamic resource change in a partition without requiring a partition reboot:

- ▶ Resource addition
- ▶ Resource removal

By achieving the resource changes sequentially in the following order on two partitions in a system, the specified resource can be moved from a partition to another partition:

1. Resource removal from a partition
2. Resource addition to another partition

This resource movement is implemented as single task on the HMC, although it is actually composed of two separate tasks on two partitions internally.

A resource is either of the following types:

- ▶ CPU

The granularity of a CPU resource of a DLPAR operation is one CPU. More than one CPU can be specified as a resource of a DLPAR operation.

A partition must be assigned at least the minimum number of processors specified in the partition profile, and it can be assigned up to the maximum number of processors specified in the partition profile.

Therefore, you can dynamically add or remove processors for that partition within the range of the minimum and maximum values.

- ▶ Memory

The granularity of a memory resource of a DLPAR operation is 256 MB.¹⁰ Multiplies of 256 MB memory can be specified as a resource of a DLPAR operation.

¹⁰ This memory region is referred to as a logical memory block (LMB).

A partition must be assigned at least the minimum size of memory specified in the partition profile, and it can be assigned up to the maximum size of memory specified in the partition profile.

Therefore, you can dynamically add or remove memory for that partition within the range of the minimum and maximum values.

► I/O resource

The granularity of an I/O resource of a DLPAR operation is a PCI slot with a PCI adapter. Multiple I/O slots can be specified as a resource of a DLPAR operation. If a PCI adapter has multiple ports, all the ports and devices configured beneath the ports are treated as a resource.

For example, if a 10/100 4-Port Ethernet adapter (FC 4961) is selected, all Ethernet devices (entX0) and interfaces (enX) configured on this adapter are treated as a single resource.

A partition must be assigned all the adapters specified as *required* in the partition profile, and it can be assigned adapters specified as *desired* in the partition profile (see “Three kinds of values for resource assignment” on page 715). Therefore, you can dynamically add or remove only adapters specified as desired for that partition.

You cannot remove I/O slots listed as required; however, you can remove I/O slots listed as desired, or those that were added as a result of a DLPAR operation. In other words, a partition can currently contain an I/O slot that is *not* listed as either desired or required in the active partition profile.

Note: A DLPAR operation can perform only one type of resource change. You cannot add and remove memory to and from the same partition in a single DLPAR operation. Also, you cannot move CPU and memory from a partition to another partition in a single DLPAR operation.

Resources removed from a partition are marked free (free resources) and owned by the global firmware of system; you can consider these resources as kept in the “free resource pool.” Free resources can be added to any partition in a system as long as the system has enough free resources.

It is imperative to understand that the DLPAR function is not solely provided by AIX 5L Version 5.2, but it is supported by the integration of following components:

► Hardware

A partitioning-capable pSeries server model is required (see Table 16-4 on page 727).

► Firmware

Depending on the models you have selected, a firmware update might be required. To download and recognize the firmware level, use the following link.

<http://techsupport.services.ibm.com/server/mdownload/>

Table 16-4 provides the minimum firmware level for dynamic logical partitioning.

Table 16-4 The minimum firmware level for dynamic logical partitioning

Official product model name	Minimum firmware level for DLPAR	M/T-MDL
IBM @server pSeries 690 Model 681	RH021114	7040-681
IBM @server pSeries 670 Model 671	RH021114	7040-671
IBM @server pSeries 655 Model 651	RJ030206	7039-651
IBM @server pSeries 650 Model 6M2	Any version	7038-6M2
IBM @server pSeries 630 Model 6C4	RR021114	7028-6C4
IBM @server pSeries 630 Model 6E4	RR021114	7028-6E4

► HMC

HMC software Release 3 Version 1 or later is required.

► Operating system

AIX 5L Version 5.2 or later is required.

If one of these components does not satisfy the requirement to support DLPAR, the function is not available. For example, if a partition is installed with AIX 5L Version 5.1, that partition does not support DLPAR, although the other partitions installed with AIX 5L Version 5.2 support DLPAR.

16.4 Managing partitions

This section describes how to use the IBM Hardware Management Console for pSeries (HMC) to manage partitioning-capable pSeries servers.

16.4.1 Hardware Management Console (HMC) user interface

The HMC provides a graphical user interface for configuring and operating single or multiple managed systems. It consists of a 32-bit Intel-based desktop PC with a DVD-RAM drive running the Linux operating system Release 7.2. The application environment, with a set of hardware management applications for configuration and partitioning, is written in Java. The applications are based on the object-oriented schema using the Common Information Model (CIM), an industry standard sponsored by the Distributed Management Task Force (DMTF). A CIM object manager acts as repository and database look-up for all managed objects.

The graphical user interface is based on the AIX 5L Version 5.2 Web-based System Manager, which allows the management integration of other HMCs or pSeries systems running AIX 5L Version 5.1 and 5.2. Except for IBM customer engineers and debugging purposes, the native Linux interfaces are hidden from the user and are not accessible. No Linux skills are required to operate the HMC. The graphical user interface can display dynamic events and static information from pSeries machines running AIX as well as from partitions on any partitioning-capable pSeries servers.

16.4.2 HMC graphical user interface

The HMC graphical user interface has the same appearance, key concepts, and basic tasks and tools as the AIX 5L Version 5.2 Web-based System Manager. For further information about the Web-based System Manager, please refer to *AIX 5L Version 5.2 Web-based System Manager Administration Guide*, available at:

<http://techsupport.services.ibm.com/server/library>

The HMC graphical user interface is composed of several elements, as shown in Figure 16-5 on page 729.

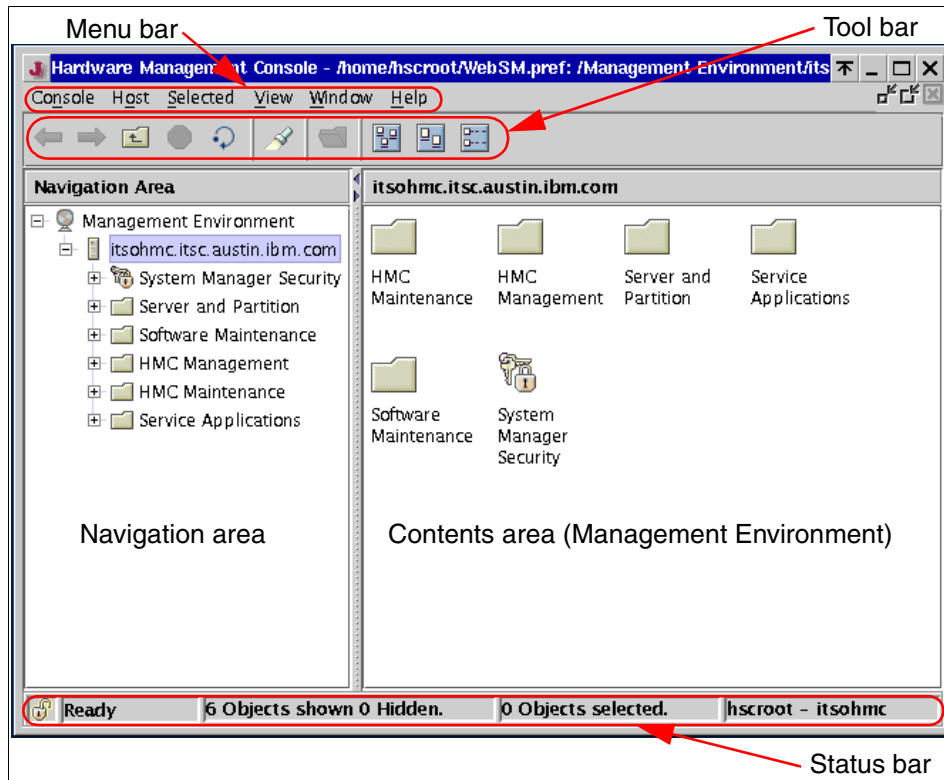


Figure 16-5 HMC graphical user interface

16.5 IBM Hardware Management Console for pSeries

In order to configure and administer a partitioning-capable pSeries server, you must attach at least one IBM Hardware Management Console for pSeries (HMC) to the system.

Note: The 7315-C03 and 7315-CR2 are the new HMCs, Feature 7315 and 7316 and also M/T 7315-C01 and 7315-C02 are withdrawn from marketing.

This section describes the Hardware Management Console (HMC) as an dedicated workstation that provides functions for the attached IBM @server pSeries server. Figure 16-6 on page 730 shows a picture of a HMC.



Figure 16-6 Picture of a desktop HMC (a rack-mount version is also available)

HMC highlights

This section describe the model highlights:

- ▶ Logical Partitioning (LPAR) control
- ▶ Clustered server control
- ▶ Capacity Upgrade on Demand (CUoD) resource control
- ▶ Creation and storage system and partition profiles
- ▶ Booting, starting, and stopping the system or individual partitions
- ▶ Displaying system and partition resources and status
- ▶ An embedded DVD-RAM, to allow easy creation and storage of configuration backup information.
- ▶ A service focal point, providing tools for problem determination and service support, such as call-home and error log notification through an analog phone line.
- ▶ Support for a 20.1-in. TFT L200p Flat Panel Monitor (FC 3636)

Description

The HMC is a dedicated desktop workstation that provides a graphical user interface for configuring and operating pSeries servers functioning in either non-partitioned or in the Full System Partition. It is configured with a set of hardware management applications for configuring and partitioning the server.

One HMC is capable of controlling multiple pSeries servers. At the time this redbook was written, a maximum of 16 non-clustered pSeries servers and a maximum of 64 partitions are supported by one HMC.

The HMC provides two native serial ports. One serial port should be used to attach a modem for the Service Agent. The second port can be used to attach a server. If multiple servers are attached to the HMC, additional serial ports are necessary. The ports can be provided by adding a maximum of two of the following features to the HMC:

- ▶ 8-Port Asynchronous Adapter (FC 2943)
- ▶ 128-Port Asynchronous Controller (FC 2944)

Note: To ensure that the Asynchronous adapter is installed in the HMC and not in the server, make sure that the adapter is configured as a feature of the HMC at the time of order.

The HMC is connected with special attachment cables to the HMC ports of the partitioning-capable pSeries server models. Only one serial connection to a server is necessary despite the number of partitions.

The following cables are available:

- ▶ Attachment Cable, HMC to host, 15 meters (FC 8121)
- ▶ Attachment Cable, HMC to host, 6 meters (FC 8120)

With these cables, the maximum length from any server to the HMC is limited to 15 meters.

To extend this distance, a number of possibilities are available:

- ▶ Another HMC can be used for remote access. This remote HMC must have a network connection to the HMC, which is connected to the servers.
- ▶ AIX 5L Web-based System Manager client can be used to connect to the HMC over the network or the Web-based System Manager PC client can be used, which runs on a Windows operating system-based or Linux operating system-based system.
- ▶ When a 128-Port Asynchronous Controller is used, the RS-422 cables connect to a remote asynchronous node (RAN) breakout box, which can be up to 330 meters. The breakout box is connected to the HMC port on the server using the attachment cable. When the 15-meter cable is used, the maximum distance the HMC can be is 345 meters, providing the entire cable length can be used.

The HMC provides a set of functions that are necessary to manage partition configurations by communicating with the service processor, as shown in Figure 16-7. These functions include:

- ▶ Creating and storing partition profiles that define the processor, memory, and I/O resources allocated to an individual partition.
- ▶ Starting, stopping, and resetting a system partition.
- ▶ Booting a partition or system by selecting a profile.
- ▶ Displaying system and partition status.

In a non-partitionable system, the LED codes are displayed in the operator panel. In a partitioned system, the operator panel shows the word LPAR instead of any partition LED codes. Therefore, all LED codes for system partitions are displayed over the HMC.

- ▶ Using a virtual console for each partition or controlled system.

With this feature, every partition can be accessed over the serial HMC connection to the server. This is a convenient feature when the partition is not reachable across the network.

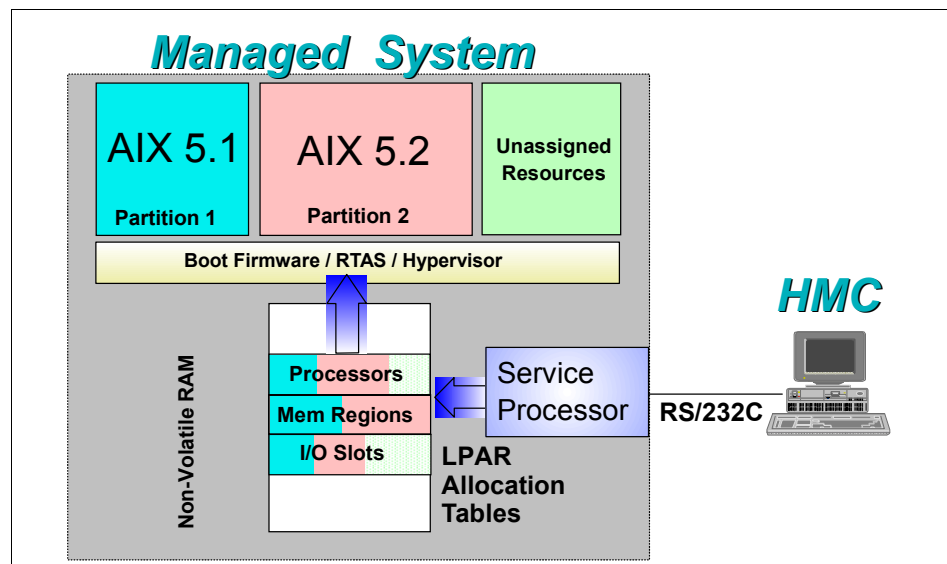


Figure 16-7 Communication between the HMC and the service processor

The HMC also provides a Service Focal Point for the systems it controls. It is connected to the service processor of the system using the dedicated serial link and must be connected to each partition using an Ethernet LAN for Service Focal Point and to coordinate dynamic logical partitioning operations. The HMC

provides tools for problem determination and service support, such as call-home and error log notification through an analog phone line.

Technical overview

The following gives you a technical overview of the HMC.

Intel Pentium-based desktop workstation with:

- ▶ 1 GB of system memory
- ▶ 40 GB minimum hard disk
- ▶ DVD-RAM for backup
- ▶ Two integrated serial ports
- ▶ One graphics port
- ▶ One integrated Ethernet port
- ▶ Six USB ports
- ▶ Three PCI slots

The following publications are shipped with the HMC. Additional copies are available.

- ▶ *System Unit Safety Information, SA23-2652*
- ▶ *Hardware Management Console for pSeries Installation and Operations Guide, SA38-0590.*

Physical specifications

This section describe the physical specifications of the HMC:

- ▶ Height: 140 mm (5.5 inches)
- ▶ Width: 425 mm (16.7 inches)
- ▶ Depth: 425 mm (16.7 inches)
- ▶ Weight: 12.0 kgs (26.5 lb.)

Operating environment

This section describe the operating environment of the HMC:

- ▶ Temperature: 10° to 35° C (50° to 95° F)
- ▶ Relative Humidity: 8% to 80%
- ▶ Sound Power: 4.4 Bels (maximum)

Power requirements and thermal output

This section describes the power requirements and the thermal output of the HMC:

- ▶ Operating voltage: 110V or 220V AC 50/60 Hz (switch control on back)
- ▶ Electrical output: 185 watts (maximum)
- ▶ Power source loading: 0.30 kva (maximum)
- ▶ 75 joules/sec (257 Btu/hr) (minimum)
- ▶ 230 joules/sec (789 Btu/hr) (maximum)

Configuration notes

The following discuss the configuration notes of the HMC.

- ▶ Each HMC must be configured with a graphics display, a keyboard, and a mouse. Existing IBM P76, P77, P260, and P275 graphics displays can be used. Various native language keyboards are available for the HMC. The HMC display menus and control information have been translated into the following languages:
 - English
 - French
 - German
 - Hungarian
 - Italian
 - Japanese
 - Russian
 - Slovakian
 - Spanish
- ▶ One HMC can control multiple pSeries servers. The number of servers each HMC can control varies by server size and complexity.
- ▶ In nonclustered environments, one HMC can control up to one of the following:
 - 12 pSeries 670 or p690 servers with up to 64 LPARS.
 - 16 pSeries 655 servers with up to 32 LPARS.
 - 16 pSeries 630 or 650 servers with up to 64 LPARS.
 - 16 pSeries servers and 64 LPARs in a mixed server environment.
 - A mixed server environment can contain a combined maximum of eight p670 or p690 servers.

- ▶ It is not possible to test each combination of servers and LPARs in a mixed server environment configuration. Performance may vary depending on the unique combination of servers and LPARs implemented.
- ▶ If redundant HMC function is desired, the pSeries servers can be attached to two separate hardware management consoles.
- ▶ An Ethernet connection between the HMC and each active partition on the pSeries server is recommended. It is required for servers operating in dynamic LPAR mode. This connection is utilized to provide:
 - Additional systems management, such as WebSM management of AIX in the individual partitions
 - Collection and passing of hardware service events to the HMC for automatic notification of error conditions to IBM
 - Total system inventory collection
- ▶ The HMC provides two integrated serial ports, an integrated 10/100/1000 Ethernet port, USB ports, and three PCI slots.
- ▶ One serial port is required for each pSeries server controlled by the HMC. An additional serial port is required for modem attachment if the Service Agent call-home function is implemented. When more than two serial ports are required, 8-port (FC 2943) and 128-port (FC 2944) asynchronous PCI adapters should be utilized.
- ▶ A combined maximum of two 8-Port and 128-Port Asynchronous Adapters are allowed per HMC.
- ▶ New Hardware Management Consoles are shipped with preloaded Hardware Management Console Software for pSeries (5639-N47) Version 1 Release 3 or later software. This software supports dynamic LPAR and Capacity Upgrade on Demand functions. To implement dynamic LPAR or CUoD function on previously shipped p670 (7040-671) or p690 (7040-681) servers, AIX 5L Version 5.2 or later systems microcode (machine code) at the 10/2002 level or later are required.

16.6 Electronic Service Agent

This section introduces general concepts of Electronic Service Agent.

16.6.1 Overview

Electronic Service Agent (also known as Service Agent) is an application program that runs on either AIX or Linux to monitor the system for hardware errors. Service Agent is an attempt to provide a proactive, predictive problem

analysis and notification by capturing diagnostic results or error log entries and sending the information to the IBM support center so that it can be analyzed and action can be taken beforehand. This not only helps reduce exposure to outages and unnecessary system downtime, but also minimizes the resources and efforts required to run and maintain the systems. Service Agent is also used by Inventory Scout to send VPD (vital product data) information to IBM or download firmware or microcode from IBM.

Figure 16-8 illustrates a typical Service Agent monitored network and how it relates to IBM.

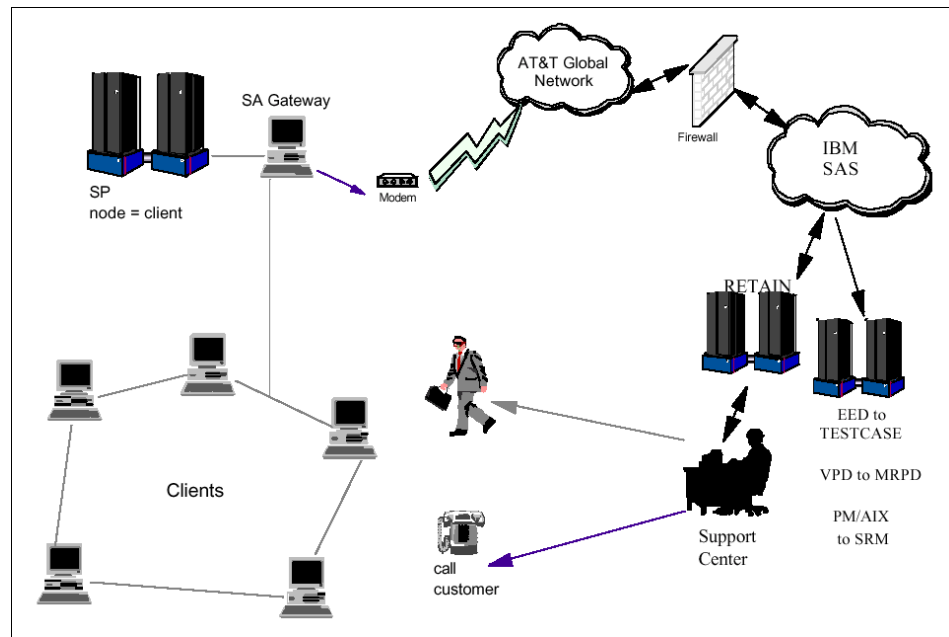


Figure 16-8 Typical Service Agent configuration

Service Agent Gateway Server is the first machine on which you installed Service Agent code. This is the machine that contains the Service Agent configuration database and has the modem connection to IBM. The Gateway machine can only be relocated to another machine by first removing the Service Agent code, then installing the Service Agent on the new Gateway machine. If you have only one machine in your system, that machine is the gateway server. If you have many machines, each of them can be a gateway server. However, if you have any number of gateway servers, you are required to have the same number of modems to connect to IBM. Additionally, you have to manage n gateway servers, since each gateway server has to be managed separately. A more practical configuration is to have one gateway server and $n-1$ Service

Agent clients. The configuration requires only one modem and you can manage all of them from that gateway server.

IBM Service Agent Server (SAS) is located at an IBM site. It is the server to which the data sent from your machines is stored, analyzed, and acted on by IBM.

Note: You must define and register the machines you wish to monitor with Service Agent to enable error detection. If you do not define and register the machines, Service Agent will not capture information.

On pSeries systems managed by the HMC, as shown in Figure 16-9, the primary path for system hardware errors detection and analysis consists of the diagnostics function provided by AIX, the service processor, and the Service Focal Point, which we will discuss in the next section.




Figure 16-9 Service Agent with the HMC

In this configuration, the Service Agent gateway process running on the HMC places service calls to the IBM support center using a dial-up connection with the attached modem, if necessary.

Figure 16-10 on page 738 shows the initial configuration panel of Service Agent on the HMC.




Figure 16-10 Service Agent on the HMC

For more information on the Service Agent, please refer to the *Hardware Management Console for pSeries Installation and Operations Guide*, SA38-0590.

16.6.2 Functions

Service Agent provides key features, including:

- ▶ Automatic problem analysis.
- ▶ Problem-definable threshold levels for error reporting.
- ▶ Automatic problem reporting; service calls placed to IBM without intervention.

- ▶ Automatic customer notification.
- ▶ Commonly viewed hardware errors. You can view hardware event logs for any monitored machine on the network from any Service Agent host user interface.
- ▶ High-availability cluster multiprocessing (HACMP) support for full fallback, includes high-availability cluster workstation (HACWS) for 9076.
- ▶ Network environment support with minimum telephone lines for modems.
- ▶ CUoD enable requirement on non HMC servers.
- ▶ VPD data can be sent to IBM.
- ▶ Software product data install and fix information will be sent to IBM.
- ▶ Using Performance Management, PM/AIX data can be automatically reported to IBM.
- ▶ Access to IBM MicroCode down load capabilities.

16.6.3 System support

This level of Service Agent supports all pSeries and RS/6000 machine types, including the 9076 (SP) or cluster configurations that have concurrent diagnostics installed and are not under control of a HMC. The application will run on HMC controlled partitions, but will not do data or error capture.

The HMC allows control of many hardware management tasks for your new pSeries servers, including configuring logical and affinity partitions. In this environment, the HMC also does hardware data capture and additional data collection normally done by SA. A different version of SA exists within the HMC, but only performs the call home communications function. AIX stand-alone SA and HMC currently cannot share the same modem.

16.7 Service Focal Point

This section provides information about using the Service Focal Point application.

16.7.1 Overview

The Service Focal Point is a system infrastructure that manages serviceable event information for the system building blocks. It includes resource managers that monitor and record information about different objects in the system. It is designed to filter and correlate events from the resource managers and initiate a

call to the service provider when appropriate. It also provides a user interface that allows a user to view the events and perform problem determination.

Traditional service strategies become more complicated in a partitioned environment. Each partition runs on its own, unaware that other partitions exist on the same system. If one partition reports an error for a shared resource, such as a managed system power supply, other active partitions report the same error. To enable service representatives to avoid long lists of repetitive call-home information, the HMC provides the Service Focal Point application. Service Focal Point recognizes that these errors repeat, and filters them into one serviceable event for the service representative to review.

Upon hardware failure events, the corresponding error entry is notified from the partition to the HMC, as shown in Figure 16-11. Also, it is worth noting that Service Focal Point only collects hardware errors, such as PERMANENT errors from AIX (marked as P) and NON BOOT errors from the service processor.




Figure 16-11 Error reporting and consolidation

From the Service Focal Point interface, you can execute maintenance procedures, such as examining the error log history, checking for components requiring replacement, and performing a Field Replaceable Unit (FRU) replacement. If Service Agent is configured on the HMC, the serviceable events are automatically sent to IBM (call-home support) for automatic generation of a maintenance request.

16.7.2 Using the Hardware Service Function

This function allows you either to just identify a frame when you have several frames connected to your HMC, or to turn off the rack indicator light. You are also able to get a Field Replaceable Unit (FRU) list when the rack indicator light is lit and check which component has problems. When a component is shown here with the LED state ON, it is much easier to identify the failing component.

To use this function, do the following:

1. Click on the + mark to the left of the Service Applications in the Navigation Area.
2. Select **Service Focal Point** in the Navigation Area. You will see the Service Focal Point task panel, as shown in Figure 16-12.
3. Select **Hardware Service Functions**.

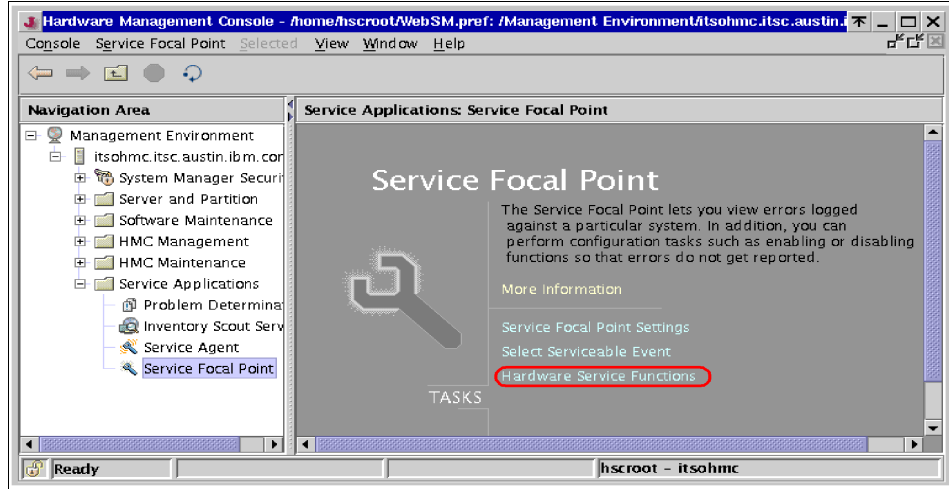


Figure 16-12 Service Focal Point: Hardware Service Functions

4. You will see the Hardware Service Management: Overview window, as shown in Figure 16-13 on page 742. Select the managed system for which you want to check the LED state, then select the **List FRUs** button.




Figure 16-13 Hardware Service Functions overview

5. You will see the FRU LED Management window, shown in Figure 16-14. If any of the LEDs are ON, it would mean that the system has a problem with the indicated component. If the Service Agent is configured to notify IBM of the errors to IBM, then IBM customer service representatives will be informed of the problem

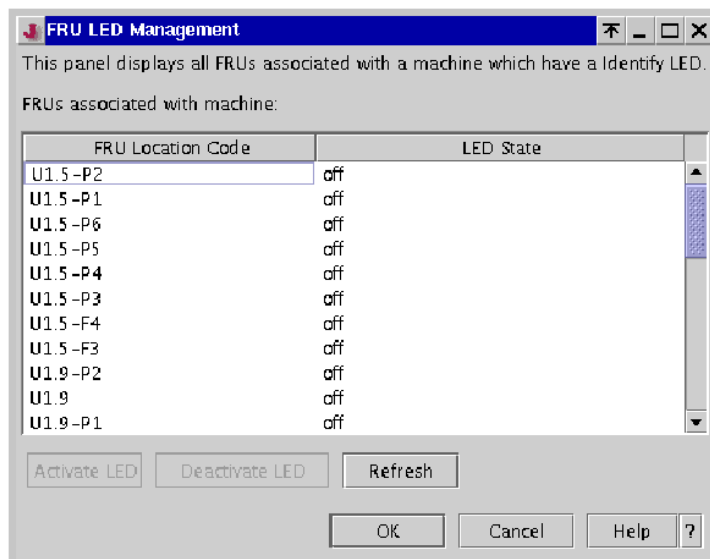


Figure 16-14 FRU LED management

For more information, please refer to *Hardware Management Console for pSeries Installation and Operations Guide*, SA38-0590.

16.8 Inventory Scout Services

In this section, we will discuss the general concepts of Inventory Scout Services.

16.8.1 Overview

Inventory Scout Services is a tool that surveys managed systems for hardware and software information. Recent releases of AIX and HMC code provide an automatic configuration mechanism and eliminate the need for you to manually reconfigure Inventory Scout. Depending on the levels of your HMC and partition software, you may be required to manually configure partitions that you create in order to perform Inventory Scout tasks. Depending on your HMC level you may also have to reconfigure Inventory Scout whenever you change partition information. You can use the HMC to perform the following Inventory Scout Services tasks:

- ▶ Inventory Scout Profile Configuration
- ▶ Conduct Microcode Survey
- ▶ Collect VPD Information
- ▶ Restart Inventory Scout Daemon

Perform each of these tasks by entering information into one or more windows. These functions will be covered in detail in the next section.

16.8.2 Functions

Inventory Scout Services are pSeries tools that perform the following two functions:

- ▶ Microcode Discovery Service

Generates a real-time comparison report showing subsystems that may need to be upgraded. This microcode survey function is similar to the microcode survey feature discussed in “Microcode Updates” on page 189, but without the updates capability. For further information about Microcode Discovery Service, visit the following URL:

<http://techsupport.services.ibm.com/server/aix.invscoutMDS>

- ▶ VPD Capture Service

Transmits your server's vital product data (VPD) information to IBM. For further information about VPD Capture Service, visit the following URL:

<http://techsupport.services.ibm.com/server/aix.invscoutVPD>

These tasks can be performed by selecting the following two tasks, as shown in Figure 16-15:

- ▶ Conduct Microcode Survey
- ▶ Collect VPD Information

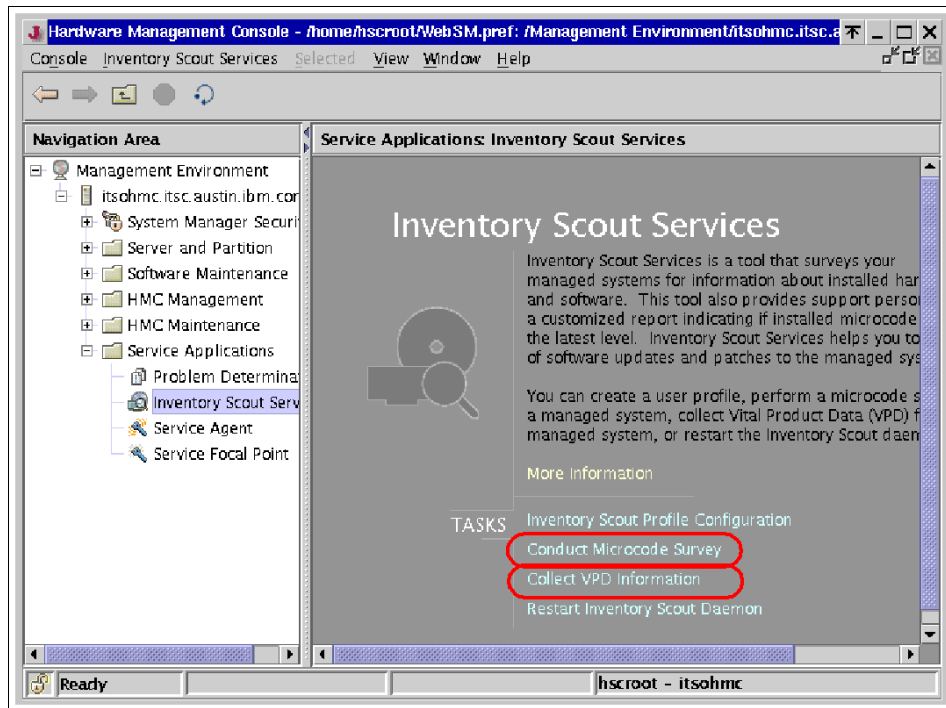


Figure 16-15 Inventory Scout Services

By selecting these two tasks on the HMC, you can check the managed system for needed microcode updates and collect vital product data (VPD) from partitions or the Full System Partition. If Service Agent is configured on the HMC, these files are automatically sent to IBM. You can optionally save files on DOS-formatted diskettes on the HMC diskette drive.

You can invoke the `invscout` command on AIX using the command line interface to Inventory Scout. This is equivalent to the Conduct Microcode Survey task on the HMC, and the command examines the microcode levels of all hardware

components visible in the partition, then accumulates the results in the microcode survey upload file.

On the pSeries 670 and pSeries 690 managed by an HMC, the accumulated file can be sent to IBM using the following steps:

1. The file is transferred to the HMC from the partition.
2. The file is sent to IBM from the HMC.

Transferring files to the HMC from partitions

If the following conditions are met, files will be automatically transferred to the HMC from the partition after the initial automatic configuration is made on the HMC:

- ▶ The firmware level of pSeries 670 or pSeries 690 is a October 2002 system microcode update or later.
- ▶ The software level of HMC is Release 3 or higher.
- ▶ The partition is installed with either of the following:
 - AIX 5L Version 5.1 with 5100-03 Recommended Maintenance Level
 - AIX 5L Version 5.2

Otherwise, Inventory Scout uses its own authentication method between the partition and the HMC in order to talk to the Inventory Scout daemon (invscoutd) on AIX. Therefore, it requires the following additional setup on AIX in a partition:

- ▶ A user, invscout, must be defined on the partition.
- ▶ A password (for example, invscout) must be set for the invscout user.



Site and hardware planning information

This appendix includes site and hardware planning information and physical planning diagrams for the following pSeries models and racks:

- ▶ 7026 Model B80
- ▶ 7028 Model 6C1 and 6E1
- ▶ 7028 Model 6C4 and 6E4
- ▶ 7029 Model 6C3 and 6E3
- ▶ 7038 @server pSeries 650
- ▶ 7039 @server pSeries 655
- ▶ 7040 @server pSeries 670
- ▶ 7040 @server pSeries 690
- ▶ 7043 43P Model 150
- ▶ 7044 44P Model 170
- ▶ 7311 Model D10
- ▶ 7311 Model D20
- ▶ 9112 Model 265
- ▶ 9114 Model 275
- ▶ 7014 Model T00 Rack
- ▶ 7014 Model T42 Rack

7026 Model B80

Table A-1 summarizes the site and hardware planning information for the 7026 Model B80.

Table A-1 7026 Model B80 site and hardware planning information

Dimensions				
Height	217 mm (8.6 in.)			
	5 EIA Units			
Width ⁴	482 mm (19 in.)			
Depth	617 mm (24.3 in.)			
Weight				
Minimum Configuration	36.5 kg		80.3 lbs	
Maximum Configuration	45.0 kg		99.3 lbs	
Electrical				
Power source loading (typical in kVA)	0.29			
Power source loading (max. in kVA)	0.46			
Voltage range (V ac)				
- US, World Trade, and Japan	100 to 127 or 200 to 240 (autoranging)			
Frequency (Hertz)	50 to 60			
Voltage range (V dc)	-48			
Thermal output (typical)	1024 BTU/hr			
Thermal output (maximum)	1536 BTU/hr			
Power requirements (typical)	300 Watts			
Power requirements (maximum)	450 Watts			
Power factor - US, World Trade, Japan	0.98			
Inrush current ³	30 amps			
Maximum altitude	2135 m (7000 ft.)			
Temperature Range	Operating		Non-Operating (Power Off)	
	10 to 40° C		10 to 52° C	
	(50 to 104° F)		(50 to 126° F)	
Humidity Requirements	Operating		Non-Operating (Power Off)	
(Noncondensing)	8 to 80%		8 to 80%	
Wet Bulb	27° C (80° F)		27° C (80° F)	
Noise Emissions¹	Operating		Idle	
LWAd	6.1 bels		5.9 bels	
LpAm	NA		NA	
<LpA>m	46 dBA		44 dBA	
Clearances	Front	Back	Left	Right

Install/Air Flow²	Maintenance of proper service clearance should allow proper air flow.
Service	See service clearances for the 7014 T00 Rack.

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.
6. For systems with FC 6120: 80/160 GB Internal Tape Drive with VXA Technology or FC 6134: 60/150 GB 16-bit 8 mm Internal Tape Drive, the maximum operating temperature is limited. For additional information about this limitation, refer to Table A-2.

Table A-2 lists maximum operating temperatures for system features at various altitudes.

Table A-2 Maximum operating temperatures

	0	305 m (1000 ft.)	610 m (2000 ft.)	914 m (3000 ft.)	1219 m (4000 ft.)	1524 m (5000 ft.)	1829 m (6000 ft.)	2134 m (7000 ft.)
FC 6120	39° C (102° F)	39° C (102° F)	39° C° (102° F)	39° C (102° F)	36° C (97° F)	35° C (95° F)	33° C (91° F)	31° C (88° F)
FC 6134	37° C (99° F)	37° C (99° F)	37° C (99° F)	37° C (99° F)	36° C (99° F)	35° C (95° F)	33° C (91° F)	31° C (88° F)

7028 Model 6C1 and 6E1

Table A-3 summarizes the site and hardware planning information for the 7028 Model 6C1 and 6E1.

Table A-3 7028 Model 6C1 and 6E1 site and hardware planning information

Dimensions	Rack (Model 6C1)	Tower (Model 6E1)
Height	215 mm (8.5 in.) 5 EIA Units	426 mm (16.8 in.)
Width	426 mm (16.8 in.)	215 mm (8.5 in.)
Depth	617 mm (24.0 in.)	617 mm (24.0 in.)
Weight		
Minimum Configuration		35.5 kg (78 lbs)
Maximum Configuration		43.1 kg (94.8 lbs)
Electrical		
Power source loading (typical in kVA)		0.40
Power source loading (max. in kVA)		0.30
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	979 BTU/hr	
Thermal output (maximum)	1306 BTU/hr	
Power requirements (typical)	288 Watts	
Power requirements (maximum)	384 Watts	
Power factor - US, World Trade/Japan	0.96	
Inrush current ²	70 Amps	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating	Non-Operating
	16 to 32° C (61 to 90° F)	10 to 43° C (50 to 109° F)
Humidity Requirements (Noncondensing)	Operating	Non-Operating
	8 to 80%	8 to 80%
Wet Bulb	27° C (80° F)	27° C (80° F)
Model 6E1 Noise Emissions	Operating	Idle
L _{WAd}	6.1 bels	6.1 bels
<L _{pA} > _m	42 dBA	41 dBA
Model 6C1 Noise Emissions	Operating	Idle
L _{WAd}	6.2 bels	5.9 bels
<L _{pA} > _m	44 dBA	41 dBA
Install/Air Flow²	Maintenance of proper service clearance should allow proper air flow.	
Service	See service clearances for the 7014 T00 Rack.	

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.

7028 Model 6C4 and 6E4

Table A-4 summarizes the site and hardware planning information for the 7028 Model 6C4 and 6E4.

Table A-4 7028 Model 6C4 and 6E4 site hardware planning information

Dimensions	Rack (Model 6C4)	Tower (Model 6E4)
Height	172.8 mm (6.8 in.) 4 EIA Units	530 mm (20.9 in.)
Width	444 mm (16.8 in.)	300 mm (8.5 in.)
Depth	609.6 mm (24.0 in.)	725 mm (28.5 in.)
Weight		
Minimum Configuration		32 kg 70.4 lbs
Maximum Configuration		47.3 kg 104.8 lbs
Electrical		
Power source loading (typical in kVA)	1-way, 2-way processors: 0.348, 4-way processor: 0.522	
Power source loading (max. in kVA)	1-way, 2-way processors: 0.522, 4-way processor: 0.783	
Voltage range (V AC)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	1-way, 2-way processors: 1129 Btu/hr, 4-way processor: 1693 Btu/hr	
Thermal output (maximum)	1-way, 2-way processors: 1693 Btu/hr, 4-way processor: 2540 Btu/hr	
Power requirements (typical)	1-way, 2-way processors: 330 watts, 4-way processor: 500 watts	
Power requirements (maximum)	1-way, 2-way processors: 500 watts, 4-way processor: 750 watts	
Power factor - US, World Trade/Japan	0.96	

Inrush current ²	50 Amps	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating 5 to 35° C (41 to 95° F)	Non-Operating 10 to 52° C (50 to 126° F)
Humidity Requirements (Noncondensing)	Operating 8 to 80%	Non-Operating 8 to 80%
Wet Bulb	27° C (80° F)	27° C (80° F)
Model 6E1 Noise Emissions	Operating	Idle
L _{WAd}	6.1 bels	6.0 bels
<L _{pA} > _m	44 dBA	43 dBA
Model 6C1 Noise Emissions	Operating	Idle
L _{WAd}	6.0 bels	5.9 bels
<L _{pA} > _m	42 dBA	41 dBA
Install/Air Flow²	Maintenance of proper service clearance should allow proper air flow.	
Service	See service clearances for the 7014 T00 Rack.	

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.
6. Levels apply to the following hardware configuration: 2-way 1 GHz processor, three Hard files, 2048 GB of RAM, Redundant system (Two 645 watt power supply, two processor fans).
7. Levels apply to the following hardware configuration: 2-way 1 GHz processor, two Hard files, 2048 GB of RAM, Non-redundant system (One 645 watt power supply, one processor fan).
8. All measurements made in accordance with ISO 7779, and declared in conformance with ISO 9296.

7029 Model 6C3 and 6E3

Table A-5 summarizes the site and hardware planning information for the 7029 Model 6C3 and 6E3.

Table A-5 7029 Model 6C3 and 6E3 site and hardware planning information

Dimensions	Rack (Model 6C3)	Tower (Model 6E4)
Height	178 mm (7.0 in.) 4EIA Units	533 mm (21.0 in.)
Width	437 mm (16.8 in.)	201 mm (8.5 in.)
Depth	508 mm (24.0 in.)	584 mm (28.5 in.)
Weight		
Minimum Configuration	35.5 kg (78.0 lbs)	35.5 kg (78.0 lbs)
Maximum Configuration	43.1 kg (94.8 lbs)	43.1 kg (94.8 lbs)
Electrical		
Power source loading (typical in kVA)	0.30	
Power source loading (max. in kVA)	0.50	
Voltage range (V AC)	100 to 127 or 200 to 240 (auto-ranging)	
Frequency (Hertz)	47 to 63	
Thermal output (typical)	1024 Btu/hr	
Thermal output (maximum)	1536 Btu/hr	
Power requirements (typical)	300 watts	
Power requirements (maximum)	450 watts	
Power factor - US, World Trade/Japan	0.95	
Inrush current ²	85 Amps (max. at <10 ms) 25 Amps (max. at 10 to 150 ms)	
Maximum altitude ³	2135 m (7000 ft.)	
Temperature Range	Operating 10 to 40° C (50 to 104° F)	Non-Operating 10 to 43° C (50 to 109° F)
Humidity Requirements (Noncondensing)	Operating 8 to 80%	Non-Operating 8 to 80%
Wet Bulb	27° C (80° F)	27° C (80° F)
Model 6E3 Noise Emissions	Operating	Idle
L _{WAd}	6.0 bels	6.0 bels
<L _{pA} > _m	42 dBA	42 dBA
Model 6C3 Noise Emissions	Operating	Idle
L _{WAd}	6.1 bels	6.1 bels
<L _{pA} > _m	44 dBA	44 dBA
Install/Air Flow²	Maintenance of service clearance will allow proper air flow.	
Service	See service clearances for the 7014 T00 Rack.	

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.
6. For systems with FC 6134: 60/150 GB 16-bit 8 mm Internal Tape Drive, the maximum operating temperature is limited. For additional information about this limitation, refer to Table A-6.

Table A-6 lists maximum operating temperatures for FC 6134: 60/150 GB 16-bit 8 mm Internal Tape Drive at various altitudes.

Table A-6 Operating temperature for FC 6134 at various altitudes.

	0	305 m (1000 ft.)	610 m (2000 ft.)	914 m (3000 ft.)	1219 m (4000 ft.)	1524 m (5000 ft.)	1829 m (6000 ft.)	2134 m (7000 ft.)
FC 6134	31° C (88° F)	31° C (88° F)	30° C (86° F)	30° C (86° F)	29° C (84° F)	29° C (84° F)	28° C (82° F)	281° C (82° F)

7038-6M2 p650

Table A-7 summarizes the site and hardware planning information for the 7038-6M2 p650.

Table A-7 7038-6M2 site and hardware planning information

Dimensions		
Height	351 mm	13.8 in.
Width	445 mm	17.5 in.
Depth	760 mm	29.9 in.
Weight		
	93 kg	205 lbs

Electrical			
Power source loading (typical)	8-way processor: 1.126 kVA		
Power source loading (max.)	8-way processor: 1.684 kVA		
Voltage range	200 to 240V ac, V dc not supported		
Frequency	50 or 60 Hz		
Thermal output (typical)	8-way processor: 3,652Btu/hr		
Thermal output (max.)	8-way processor: 5,461 Btu/hr		
Power requirements (typical)	8-way processor: 1,070 watts		
Power requirements (max.)	8-way processor: 1,600 watts		
Power factor	0.95		
Inrush current (2)	67 amps at 200V ac, 60 Hz 87 amps at 230V ac, 50 Hz		
Maximum altitude (3, 4)	3048 m (10000 ft.)		
Temperature Requirements (3)	Operating	Non-Operating	Storage
	10 to 38° C (50 to 100° F)	1 to 43° C (34 to 109° F)	1 to 60° C (34 to 140° F)
Humidity Requirements (4) (Noncondensing)	Operating	Non-Operating	Storage
Wet Bulb	8 to 80% 23° C (73° F)	8 to 80% 27° C (81° F)	5 to 80% 29° C (84.2° F)
Noise Emissions (1, 5, and 6)	Operating	Idle	
L _{WAd}	61 bels (5)	61 bels (5)	
<L _{pA} > _m	44 dBA (6)	44 dBA (6)	
Install/Air Flow	Maintenance of service clearance will allow proper air flow.		
Service Clearances	See “T00 and T42 Service Clearances and Caster Location” on page 31 for T00 or T42 rack service clearances.		

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. The LWAd emission increases to 6.5 bels with a configuration of one 7038-6M2 and four 7311-D10 drawers.
6. The LpA >m emission increases to 48 dBA with a configuration of one 7038-6M2 and four 7311-D10 drawers.

7039 pSeries 655

Table A-8 summarizes the components for the 7039 pSeries 655.

Table A-8 7039 pSeries 655 multiple components

Model	Description	Minimum per system	Maximum per system
7040-W42	Base Frame (Redundant power supplies as feature codes)	1	1
FC 6076	Slimline Front Door (2)	1 (1)	1 (1)
FC 6119	Acoustic Front Door (2)	1 (1)	1 (1)
FC 6078	Slimline Rear Door (2)	1 (1)	1 (1)
FC 6079	Acoustical Rear Door (2)	1 (1)	1 (1)
FC 6200 or FC 6201	Optional Integrated Battery Feature (IBF)	0	6
7039-651	Server Node (up to 8 processors, 4 GB to 32 GB memory)	1	16
7035-C01	Hardware Management Console (HMC)	0	2
7040-61D	IO Subsystem (20 PCI cards maximum, 16 DASD maximum)	0	5

Notes:

1. Either slimline doors or acoustical doors must be selected by the customer during the order process.
2. Door options determine which doors are included with your pSeries 655.

Table A-9 summarizes the site and hardware planning information for the 7039 pSeries 655

Table A-9 7039 pSeries 655 site and hardware planning information

Dimensions	Slimline Doors	Acoustical Doors
Height	2028 mm (79.84 in.)	2028 mm (79.84 in.)
Width	785 mm (30.91 in.)	785 mm (30.91 in.)
Depth	1443 mm (56.81 in.)	1799 mm (70.83 in.)
Weight		
Weight without IBF (max. config.)	1487 kg (3279 lbs.)	1496 kg (3299 lbs.)
Weight with IBF (max. config.)	1629 kg (3592 lbs.)	1638 kg (3612 lbs.)

Electrical			
Rated Voltage (V ac, 3 phase)	200 to 240	380 to 415	480
Rated Current, Line Cord with 60-A Plug, FC 8688 or 8689 (amps, per phase)	48		
Rated Current, All Other Line Cords (amps, per phase)	60	32	24
Frequency (Hertz)	50 to 60	50 to 60	50 to 60
Power requirements (max. in KW)	31.8	31.8	31.8
Power Factor US, World Trade, Japan (pf)	0.99	0.97	0.95
Thermal output (max. kBtu/hr))	102	102	102
Inrush current (Max A)	162	162	162
Maximum altitude (3, 4)	3048 m (10000 ft.)		
Temperature Requirements	Operating 10 to 32° C (50 to 90° F)	Non-Operating 10 to 43° C (50 to 109° F)	Storage 1 to 60° C (34 to 140° F)
Humidity Requirements (Noncondensing)	Operating 8 to 80%	Non-Operating 8 to 80%	Storage 5 to 80%
Wet Bulb	23° C (73° F)	27° C (81° F)	29° C (84.2° F)
Noise Emissions (non-acoustical doors)	Operating	Idle	
One processor node L_{WAd}	7.4 bels	7.4 bels	
One processor node $\langle L_{pA} \rangle_m$	57 dBA	57 dBA	
Three processor nodes L_{WAd}	8.2 bels	8.2 bels	
Three processor nodes $\langle L_{pA} \rangle_m$	64 dBA	64 dBA	
Sixteen processor nodes L_{WAd}	8.7 bels	8.7 bels	
Sixteen processor nodes $\langle L_{pA} \rangle_m$	69 dBA	69 dBA	
Noise Emissions (acoustical doors)	Operating	Idle	
One processor node L_{WAd}	6.7 bels	6.7 bels	
One processor node $\langle L_{pA} \rangle_m$	50 dBA	50 dBA	
Three processor nodes L_{WAd}	7.5 bels	7.5 bels	
Three processor nodes $\langle L_{pA} \rangle_m$	57 dBA	57 dBA	
Sixteen processor nodes L_{WAd}	8.0 bels	8.0 bels	
Sixteen processor nodes $\langle L_{pA} \rangle_m$	62 dBA	62 dBA	
Install/Air Flow	Maintenance of service clearance will allow proper air flow.		
Service Clearances	See note 6.		

Notes:

1. Doors are not installed during product shipment to the customer.

2. Refer to the table at page 92 in the *RS/6000 and IBM @server pSeriesSite and Hardware Planning Information*, SA38-0508 manual for the approximate weight of your system configuration.
3. Inrush currents occur only at initial application of power (very short duration for charging capacitors). No inrush currents occur during the normal power off-on cycle.
4. When an IBM-approved vapor bag and desiccant packets are used to protect the system, storage specifications are valid for six months and shipping specifications are valid for one month. Otherwise, storage and shipping specifications are valid for two weeks each.
5. The upper limit of the dry bulb temperature must be derated 1° C per 219 m (719 ft.) above 1295 m (4250 ft.). Maximum altitude for is 3048 m (10,000 ft.).
6. Service clearance is independent from weight distribution distance and must be at least 1194 mm (47 in.) for the front of the frame and 915 mm (36 in.) for the rear of the frame (measured from the base frame).
7. Weight-distribution areas should not be overlapped.
8. Floor-loading weight distribution distances should not exceed 762 mm (30 in.) in any direction when measured from the base frame.
9. Consult a professional, such as a structural engineer, if you are unsure of the floor-load rating of your facility.

Figure A-1 on page 759 shows a proposed floor layout for multiple 7039 systems.




Figure A-1 Proposed floor layout for multiple 7039 systems

Figure A-2 shows the minimum service clearance for 7039 systems with thin doors.




Figure A-2 Service clearance for 7039 systems with thin doors

Figure A-3 on page 761 shows the minimum service clearance for 7039 systems with acoustical doors.




Figure A-3 Service clearance for 7039 systems with acoustical doors

7040 pSeries 670

The 7040 pSeries 670 system consists of multiple components, as summarized in Table A-10.

Table A-10 7040 pSeries 670 multiple components

Model	Description	Minimum per system	Maximum per system
7040-61R	Base Frame (Redundant power supplies as feature codes)	1	1
FC 6070	Base Frame Universal Front Door	1	1
FC 6074	Base Frame Slimline Rear Door (2)	1 (1)	1 (1)
FC 6075	Base Frame Acoustical Rear Door (2)	1 (1)	1 (1)
FC 6200 or FC 6201	Optional Integrated Battery Feature (IBF)	0	2
7040-671	Managed Server (up to 16 processors, 4 GB to 128 GB memory)	1	1
FC 7315	Hardware Management Console (HMC)	0	2
7040-61D	IO Subsystem (20 PCI cards maximum, 16 DASD maximum)	1	3

Notes:

1. Either slimline doors or acoustical doors must be selected by the customer during the order process. Slimline doors will not meet acoustic limits for Category 1A.
2. Door options determine which doors are included with your pSeries 670.

Table A-11 summarizes the site and hardware planning information for the 7040 pSeries 670

Table A-11 7039 pSeries 670 site and hardware planning information

Dimensions	Slimline Doors	Acoustical Doors
Height	2025 mm (79.72 in.)	2025 mm (79.72 in.)
Width	785 mm (30.91 in.)	785 mm (30.91 in.)
Depth	1342 mm (52.83 in.)	1494 mm (58.83 in.)
Weight		
Weight (maximum configuration)	1085 kg (2392 lbs.)	1099 kg (2422 lbs.)

Electrical/Thermal Characteristics (3-Phase)			
Rated Voltage (V ac, 3 phase)	200 to 240	380 to 415	480
Rated Current (A, per phase)	45	25	20
Frequency (Hertz)	50 to 60	50 to 60	50 to 60
Power requirements (max. in KW)	6.7	6.7	6.7
Typical, full load Power Factor (pf)	0.99	0.97	0.93
Thermal output (max. kBtu/hr)	22.8	22.8	22.8
Inrush current (Max A)	162	162	162
Electrical/Thermal Characteristics (1-Phase)			
Rated Voltage (V ac, 1 phase)	200 to 240	380 to 415	
Rated Current (A, per phase)	38	19.5	
Frequency (Hertz)	50 to 60	50 to 60	
Power requirements (max. in KW)	6.7	6.7	
Typical, full load Power Factor (pf)	0.99	0.97	
Thermal output (max. kBtu/hr)	22.8	22.8	
Inrush current (Max A)	162	162	
Max. altitude 1.1 GHz Modules	3048 m (10000 ft.)		
Max. altitude 1.3 GHz Modules	2134 m (7000 ft.)		
Temperature Requirements	Operating 10 to 32° C (50 to 90° F) Max. of 24° C (75.2° F) with 4 mm tape or DVD RAM in rear positions of the Media Subsystem	Non-Operating 10 to 43° C (50 to 109° F)	Storage 1 to 60° C (34 to 140° F)
Humidity Requirements (Noncondensing) Wet Bulb	Operating 8 to 80% 23° C (73° F)	Non-Operating 8 to 80% 27° C (81° F)	Storage 5 to 80% 29° C (84° F)
Noise Emissions (acoustical doors)	Operating	Idle	
L _{WAd}	7.5 bels	7.5 bels	
<L _{pA} > _m	57 dBA	57 dBA	
Noise Emissions (non-acoustical doors)	Operating	Idle	
L _{WAd}	7.9 bels	7.9 bels	
L _{pA} > _m	62 dBA	62 dBA	
Install/Air Flow	Maintenance of service clearance will allow proper air flow.		
Service Clearances	See note 5.		

Notes:

1. When moving or relocating certain configurations of the system, the Bulk Power Regulators (BPRs) must be removed from the top of the rack (front and rear) to ensure product stability. Specifically, removal of the BPR from the front and rear is required in systems that have a single I/O drawer.
2. Inrush currents occur only at initial application of power (very short duration for charging capacitors). No inrush currents occur during the normal power off-on cycle.
3. When an IBM-approved vapor bag and desiccant packets are used to protect the system, storage specifications are valid for six months and shipping specifications are valid for one month. Otherwise, storage and shipping specifications are valid for two weeks each.
4. The upper limit of the dry bulb temperature must be derated 1° C per 189 m (619 ft.) above 1295 m (4250 ft.). Maximum altitude is 3048 m (10,000 ft.)
5. Service clearance is independent from weight distribution distance and must be at least 1143 mm (45 in.) for the front of the frame and 915 mm (36 in.) for the rear of the frame (measured from the base frame).
6. Weight-distribution areas should not be overlapped.
7. Floor-loading weight distribution distances should not exceed 762 mm (30 in.) in any direction when measured from the base frame.
8. Consult a professional, such as a structural engineer, if you are unsure of the floor-load rating of your facility.

Figure A-4 on page 765 shows a proposed floor layout for multiple 7040 systems.




Figure A-4 Proposed floor layout for multiple systems

Figure A-5 shows the minimum service clearance for 7040 single-frame systems with slimline doors.




Figure A-5 Service clearance for 7040 single-frame systems with slimline doors

Figure A-6 on page 767 shows the minimum service clearance for 7040 single-frame systems with acoustical doors.




Figure A-6 Service clearance for 7040 single-frame systems w/ acoustical doors

7040 pSeries 690

The 7040 pSeries 690 system consists of multiple components, as summarized in Table A-12.

Table A-12 7040 pSeries 690 multiple components

Model	Description	Minimum per system	Maximum per system
7040-61R	Base Frame (Redundant power supplies as feature codes)	1	1
FC 8691	Optional Expansion Frame	0	1
FC 6070	Base Frame Universal Front Door	1	1
FC 6071	Expansion Frame Universal Front Door	0	1
FC 6074	Base/Exp. Frame Slimline Rear Door (2)	1 (1)	2 (1)
FC 6075	Base/Exp. Frame Acoustical Rear Door (2)	1 (1)	2 (1)
FC 6200 or FC 6201	Optional Integrated Battery Feature (IBF)	0	6FC
7040-671	Managed Server (up to 16 processors, 4 GB to 128 GB memory)	1	1
FC 7315	Hardware Management Console (HMC)	0	2
FC 8692	Media Subsystem (Operation panel, 3.5-in. floppy drive, optional media devices)	1	1
7040-61D	IO Subsystem (20 PCI cards maximum, 16 DASD maximum)	1	8

Notes:

1. Either slimline doors or acoustical doors must be selected by the customer during the order process. Slimline doors will not meet acoustic limits for Category 1A.
2. Door options determine which doors are included with your pSeries 690.

Table A-13 on page 769 summarizes the site and hardware planning information for the 7040 pSeries 690.

Table A-13 7039 pSeries 690 site and hardware planning information

Dimensions	Slimline Doors	Acoustical Doors	
Height (maximum)	2025 mm (79.72 in.)	2025 mm (79.72 in.)	
Width 1 Frame	785 mm (30.91 in.)	785 mm (30.91 in.)	
Width 2 Frame	1575 mm (62.00 in.)	1575 mm (62.00 in.)	
Depth (maximum)	1342 mm (52.83 in.)	1494 mm (58.83 in.)	
Weight			
Weight 1 Frame (max. config.)	1170 kg (2580 lbs.)	1184 kg (2610 lbs.)	
Weight 2 Frames (max. config.)	1973 kg (4340 lbs.)	2000 kg (4409 lbs.)	
Electrical/Thermal Characteristics			
Rated Voltage (V ac, 3 phase)	200 to 240	380 to 415	480
Rated Current (A, per phase)	45	25	20
Frequency (Hertz)	50 to 60	50 to 60	50 to 60
Power requirements (max. in KW)	15.7	15.7	15.7
Typical, full load Power Factor (pf)	0.99	0.97	0.93
Thermal output (max. kBtu/hr)	53.3	53.3	53.3
Inrush current (Max A)	162	162	162
Maximum altitude (3, 4)	3048 m (10000 ft.)		
Temperature Requirements	Operating	Non-Operating	Storage
	10 to 32° C (50 to 90° F) Max. of 24 ° C (75.2 ° F) with 4 mm tape or DVD RAM in rear positions of the Media Subsystem	10 to 43° C (50 to 109° F)	1 to 60° C (34 to 140° F)
Humidity Requirements (Noncondensing) Wet Bulb	Operating 8 to 80% 23° C (73° F)	Non-Operating 8 to 80% 27° C (81° F)	Storage 5 to 80% 29° C (84° F)
Noise Emissions (acoustical doors)	Operating	Idle	
L _{WAd}	7.5 bels	7.5 bels	
<L _{pA} > _m	57 dBA	57 dBA	
Noise Emissions (non-acoustical doors)	Operating	Idle	
L _{WAd}	7.9 bels	7.9 bels	
L _{pA} > _m	62 dBA	62 dBA	
Install/Air Flow	Maintenance of service clearance will allow proper air flow.		
Service Clearances	See note 6.		

Notes:

1. Doors are not installed during product shipment to the customer. A maximum configured system with batteries may exceed 1134 kg (2500 lbs.).
2. When moving or relocating certain configurations of the system, the Bulk Power Regulators (BPRs) must be removed from the top of the rack (front and rear) to ensure product stability. Specifically, removal of BPRs from frame A and B in the front and rear is required in systems that have a single I/O drawer, and more than 2 BPRs installed per BPA in the primary rack.
3. Inrush currents occur only at initial application of power (very short duration for charging capacitors). No inrush currents occur during the normal power off-on cycle.
4. When an IBM-approved vapor bag and desiccant packets are used to protect the system, storage specifications are valid for six months and shipping specifications are valid for one month. Otherwise, storage and shipping specifications are valid for two weeks each.
5. The upper limit of the dry bulb temperature must be derated 1° C per 189 m (619 ft.) above 1295 m (4250 ft.).
6. Service clearance is independent from weight distribution distance and must be at least 1143 mm (45 in.) for the front of the frame and 915 mm (36 in.) for the rear of the frame (measured from the base frame).
7. Weight-distribution areas should not be overlapped.
8. Floor-loading weight distribution distances should not exceed 762 mm (30 in.) in any direction when measured from the base frame.
9. Consult a professional, such as a structural engineer, if you are unsure of the floor-load rating of your facility.

Figure A-7 on page 771 shows a proposed floor layout for multiple 7040 systems.




Figure A-7 Proposed floor layout for multiple systems

Figure A-8 shows the minimum service clearance for 7040 single-frame and double-frame systems with slimline doors.




Figure A-8 Service clearance for 7040 single/double frame with slimline doors

Figure A-9 on page 773 shows the minimum service clearance for 7040 single-frame and double-frame systems with acoustical doors.




Figure A-9 Clearance for 7040 single and double frame with acoustical doors

7043 43P Model 150

Table A-14 summarizes the site and hardware planning information for the 7043 43P Model 150.

Table A-14 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 120V AC and at 240V AC			
Maximum altitude ³	2135 m (7000 ft.)			
Temperature Range	Operating	Non-Operating		
	16 to 32° C (60 to 90° F)	10 to 43° C (50 to 110° F)		
Humidity (Noncondensing)	Operating	Non-Operating		
	8 to 80%	8 to 80%		
Wet Bulb Requirements	23° C (73° F)	27° C (80° F)		
Noise Emissions ¹	Operating	Idle		
L _{WAd}	5.4 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 Flow ²	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.			

Notes:

1. See "Noise emission notes" on page 783 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 A-10 on page 775 shows the physical planning diagram for the 7043 43P Model 150.




Figure A-10 7043 43P Model 150 physical planning diagram

7044 44P Model 170

Table A-15 summarizes the site and hardware planning information for the 7044 44P Model 170.

Table A-15 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)	0.40	
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 120V AC and at 240V ac	

Maximum altitude	2135 m (7000 ft.)			
Temperature Range	Operating		Non-Operating (Power Off)	
	16 to 32° C (60 to 90° F)		10 to 43° C (50 to 110° F)	
Humidity Requirements (Noncondensing)	Operating		Non-Operating (Power Off)	
	8 to 80%		8 to 80%	
Wet Bulb Requirements	23° C (73° F)		27° C (80° F)	
Noise Emissions¹	Operating		Idle	
L _{WAd}	5.5 bels		5.4 bels	
L _{pAm}	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.			

Notes:

1. See “Noise emission notes” on page 783 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.
4. Width measurement with feet extended.

Figure A-11 shows the physical planning diagram for the 7044 44P Model 170.




Figure A-11 7044 44P Model 170 physical planning diagram

7311 Model D10

Table A-16 summarizes the site and hardware planning information for the 7311 Model D10.

Table A-16 7311 Model D10 site and hardware planning information

Dimensions	7311-D10	Two 7311-D10s with Enclosure	
Height	170 mm (6.6 in.)	178 mm (7.0 in.)	
Width ⁴	220 mm (8.7 in.)	445 mm (17.5 in.)	
Depth	711 mm (28.0 in.)	711 mm (28.0 in.)	
Weight			
7311-D10 (max. configuration)	16.8 kg (37 lbs)	29.1 kg (86 lbs)	
Electrical			
Power source loading per 7311-D10	0.21 kVA		
Voltage range	200 to 240V ac, V dc not supported		
Frequency (Hertz)	50 to 60		
Thermal output per 7311-D10	461 BTU/hr		
Thermal output per 7311-D10 (max.)	683 BTU/hr		
Power requirements (typical)	135 Watts (per 7311-D10)		
Power requirements (maximum)	200 Watts (per 7311-D10)		
Power factor	0.91		
Inrush current per 7311-D10	64 amps		
Maximum altitude	3048 m (10000 ft.)		
Temperature Range	Operating	Non-Operating	Storage
	10 to 38° C (50 to 100° F)	1 to 60° C (34 to 140° F)	1 to 60° C (34 to 140° F)
Humidity Requirements (Noncondensing)	Operating	Non-Operating	Storage
	8 to 80%	8 to 80%	8 to 80%
Wet Bulb	23° C (80° F)	27° C (80° F)	29° C (80° F)
Noise Emissions	Operating	Idle	
LWAd one 7311-D10	5.6 bels	5.6 bels	
LWAd two 7311-D10	5.9 bels	5.9 bels	
LWAd four 7311-D10	6.2 bels	6.2 bels	
<LpA>m one 7311-D10	40 dBA	40 dBA	
<LpA>m one 7311-D10	43 dBA	43 dBA	
<LpA>m one 7311-D10	46 dBA	46 dBA	
Clearances	See Table A-20 on page 782 for the T00 or T42 rack service clearances.		
Install/Air Flow²	Maintenance of service clearance should allow proper air flow.		

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).

7311 Model D20

Table A-17 summarizes the site and hardware planning information for the 7311 Model D20.

Table A-17 7311 Model D20 site and hardware planning information

Dimensions			
Height	178 mm (7.0 in.)		
Width ⁴	445 mm (17.5 in.)		
Depth	610 mm (24.0 in.)		
Weight			
maximum configuration	45.9 kg	101 lbs	
Electrical			
Power source loading (maximum)	0.358 kVA		
Voltage range	200 to 240V ac, V dc not supported		
Frequency (Hertz)	50 to 60		
Thermal output (typical)	774 BTU/hr		
Thermal output (maximum)	1161 BTU/hr		
Power requirements (typical)	227 Watts		
Power requirements (maximum)	340 Watts		
Power factor	0.91		
Inrush current per 7311-D10	60 amps		
Maximum altitude	3048 m (10000 ft.)		
Temperature Range	Operating	Non-Operating	Storage
	5 to 35° C (41 to 96° F)	1 to 60° C (34 to 140° F)	1 to 60° C (34 to 140° F)
Humidity Requirements (Noncondensing)	Operating	Non-Operating	Storage
	8 to 80%	8 to 80%	8 to 80%
Wet Bulb	23° C (80° F)	27° C (80° F)	29° C (80° F)

Noise Emissions	Operating	Idle
LWAd	6.1 bels	6.0 bels
<LpA>m	44 dBA	43 dBA
Clearances	See Table A-20 on page 782 for the T00 or T42 rack service clearances.	
Install/Air Flow²	Maintenance of service clearance should allow proper air flow.	

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).

9112 Model 265

Table A-18 summarizes the site and hardware planning information for the 9112 Model 265.

Table A-18 9112 Model 265 site and hardware planning information

Dimensions		
Height	426 mm (16.8 in.)	
Width ⁴	215 mm (8.5 in.)	
Depth	617 mm (24.3 in.)	
Weight		
Minimum Configuration	35.5 kg	78 lbs
Maximum Configuration	43.1 kg	94.8 lbs
Electrical		
Power source loading (typical in kVA)	0.30	
Power source loading (max. in kVA)	0.40	
Voltage range (V ac)	100 to 127 or 200 to 240 (autoranging)	
Frequency (Hertz)	50 to 60	
Thermal output (typical)	979 BTU/hr	
Thermal output (maximum)	1306 BTU/hr	
Power requirements (typical)	288 Watts	
Power requirements (maximum)	384 Watts	

Power factor - US, World Trade, Japan	0.96	
Inrush current ³	70 amps	
Maximum altitude	2135 m (7000 ft.)	
Temperature Range	Operating 16 to 32° C (61 to 90° F)	Non-Operating (Power Off) 10 to 43° C (50 to 109° F)
Humidity Requirements (Noncondensing)	Operating 8 to 80%	Non-Operating (Power Off) 8 to 80%
Wet Bulb	27° C (80° F)	27° C (80° F)
Noise Emissions¹	Operating	Idle
LWAd	6.1 bels	6.1 bels
LpAm	42 dBA	41 dBA
Install/Air Flow²	Maintenance of proper service clearance should allow proper air flow.	
Service	See service clearances for the 7014 T00 Rack	

Notes:

1. See “Noise emission notes” on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.

9114 Model 275

Table A-19 summarizes the site and hardware planning information for the 9114 Model 275.

Table A-19 9114 Model 275 site and hardware planning information

Dimensions	
Height	535.9 mm (21.1 in.)
Width	190 mm (7.5 in.)
Depth	685 mm (27.0 in.)

Weight		
Minimum Configuration	32.0 kg	70.5 lbs
Maximum Configuration	43.1 kg	94.8 lbs
Electrical		
Power source loading (typical in kVA)	0.30	
Power source loading (max. in kVA)	0.50	
Voltage range (V ac)	100 to 127 or 200 to 240 (auto-ranging)	
Frequency (Hertz)	47 to 63	
Thermal output (typical)	1024 BTU/hr	
Thermal output (maximum)	1587 BTU/hr	
Power requirements (typical)	300 Watts	
Power requirements (maximum)	465 Watts	
Power factor - US, World Trade, Japan	0.95	
Inrush current ²	85 amps	
Maximum altitude ^{3,4}	2135 m (7000 ft.)	
Temperature Range^{3,6}	Operating 10 to 32° C (50 to 90° F)	Non-Operating (Power Off) 10 to 43° C (50 to 109° F)
Humidity Requirements⁴ (Noncondensing)	Operating 8 to 80%	Non-Operating (Power Off) 8 to 80%
Wet Bulb	27° C (80° F)	27° C (80° F)
Noise Emissions¹	Operating	Idle
LWAd	5.5 bels	5.5 bels
Install/Air Flow	Maintenance of proper service clearance should allow proper air flow.	
Service	See service clearances for the 7014 T00 Rack.	

Notes:

1. See "Noise emission notes" on page 783 for definitions of noise emissions positions. See noise emissions note 4.
2. Inrush currents occur only at initial application of power. No inrush occurs during normal power off-on cycle.
3. The upper limit of the dry bulb temperature must be derated 1° C per 137 m (450 ft.) above 915 m (3000 ft.).
4. The upper limit of the wet bulb temperature must be derated 1° C per 274 m (900 ft.) above 305 m (1000 ft.).
5. Levels are for a single system installed in a T00 32 EIA rack with the center of the unit approximately 1500 mm (59 in.) off the floor.

6. For systems with FC 6134: 60/150 GB 16-bit 8 mm Internal Tape Drive, the maximum operating temperature is limited.

7014 Model T00 and T42 Rack

Table A-20 summarizes the site and hardware planning information for the Model T00 Rack, while Table A-21 on page 783 refers to the Model T42 rack.

Table A-20 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.
Weigh				
Base Rack	244 kg			535 lbs
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 240V 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.)

Notes:

1. Configuration-dependent (base weight plus weight of drawers).
2. The total rack power should be derived from the sum of the power used by the drawers in the rack.

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).
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.

Table A-21 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.
Weigh		
Base Rack	261 kg	575 lbs
Full Rack	930 kg	2045 lbs
Service Clearance	Recommended minimum vertical service clearance from floor is 2439 mm or 8 feet.	
All other Specifications	For all other technical information, see Table A-20 on page 782.	

Noise emission notes

The following are noise emission notes that were referred to in the tables in this appendix:

1. LWAd is the declared sound power emission level for a production series of machines.
2. LpAm is the mean value of the sound pressure emission levels at the operator position (if any) for a production series of machines.
3. <L pA > 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.



Adapter placement guidelines

This appendix describes the adapter placement rules for the following systems. The information presented here is a subset of *RS/6000 and IBM @server pSeries Adapter Placement Reference for AIX, SA38-0538*. The most current information exists in the reference.

- ▶ 7025 pSeries 620 Model 6F0/6F1
- ▶ 7026 pSeries 640 Model B80
- ▶ 7028 pSeries 610 Model 6C1/6E1
- ▶ 7028 pSeries 630 Model 6C4/6E4
- ▶ 7029 pSeries 615 Model 6C3/6E3
- ▶ 7040 pSeries 670 Model 61D
- ▶ 7040 pSeries 690 Model 61D
- ▶ 7038 pSeries 650 Model 6M2
- ▶ 7039 pSeries 655 Model 651
- ▶ 7043 Model 150
- ▶ 7044 Model 170
- ▶ 7044 Model 270
- ▶ 7046 Model B50
- ▶ 7311 Model D10
- ▶ 7311 Model D20
- ▶ 9076 RS/6000 SP Systems
- ▶ 9112 Model 265

► 9114 Model 275

This appendix presents a brief overview of some factors and issues related to Peripheral Component Interconnect (PCI) bus adapter placement. The user can gain a better understanding of system configurations, adapter placement, and performance issues by using the reference. Each of the following sections provide system-specific slot placement information.

You can install PCI adapters with the power on in some systems. These adapters are referred to as hot-pluggable PCI adapters. Do not hot-plug any PCI adapter supporting the system's boot device residing console.

Refer to your system unit documentation to determine if your system unit supports hot-plugging adapters. If an adapter is listed as a *Type **, this adapter has not been assigned an adapter type.

Introduction

The following is a general discussion of the various considerations and restrictions regarding adapter placement and performance.

System performance

This section provides performance information related to PCI adapter placement. Understand that maximizing system performance is relative to software and hardware. Information in this section may become obsolete as new products are announced. The publication *RS/6000 and IBM eServer pSeries Adapter Placement Reference for AIX, SA38-0538* should be consulted whenever the latest information is required.

Integrated adapters

The main processor board now integrates a number of devices, 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. Integrated PCI adapters include SCSI adapters and Ethernet adapters.

32-bit versus 64-bit PCI slots

Choosing between 32-bit and 64-bit slots influences slot placement and affects performance. Higher-speed adapters use 64-bit slots because they can transfer 64 bits of data for each data transfer phase. 32-bit adapters can typically function in 64-bit PCI slots; however, 32-bit adapters still operate in 32-bit mode and offer no performance advantage in a 64-bit slot. Likewise, most 64-bit adapters can operate in 32-bit PCI slots, but the 64-bit adapter operates in 32-bit mode and reduces performance potential.

33 MHz versus 50/66 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 you plug a 33 MHz adapter 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. For systems with hot-plug PCI slots, adapters are not affected by the clock rate of other adapters, because each has its own PCI bus.

The following adapters run at 50 or 66 MHz when placed on a 50 or 66 MHz PCI bus, or 33 MHz when placed on a 33 MHz PCI bus. If you place a 33 Mhz adapter on the same 50/66 MHz PCI bus with any of these adapters, the bus will

run in 33 MHz mode, which will reduce the performance potential of the following 50/66 MHz adapters:

- ▶ FC 2969: Gigabit Ethernet Fibre, 1000 BaseT, 64-bit, 3.3/5V (Type 9-U)
- ▶ FC 2851: POWER GXT250P (Type 1-M)
- ▶ FC 2852: POWER GXT255P (Type 1-N)
- ▶ FC 2841: POWER GXT300P (Type 1-U)
- ▶ FC 2823: POWER GXT2000P (Type 1-S)
- ▶ FC 2825: POWER GXT3000P (Type 1-R)
- ▶ FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)
- ▶ FC 6228: 2 Gigabit FC Adapter 32/64-bit, 3.3/5V PCI Bus (Type 4-W)
- ▶ FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
- ▶ FC 6205: PCI Dual Channel Ultra2 SCSI (Type 4-R)
- ▶ FC 2498: PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)

Connectivity versus performance overview

You must consider some performance trade-offs when configuring your system. Installing the maximum number of adapters might affect system performance. The following paragraphs provide an overview of these considerations and how they are documented in later sections of this book.

Connectivity limits define how many specified adapters can be physically plugged into a system. This limit defines how many adapters the software and hardware can support. Some adapters have specific placement guidelines. Connectivity limits define the maximum number of adapters for connecting to networks or disks. In many cases, a disk or network has a low duty-cycle and the system needs additional adapters to retain the physical connection to all resources. In these cases, you should follow the connectivity limits.

This section also provides suggested performance limits, established to determine how many concurrently running adapters can provide good performance. As you add adapters (with each adapter performing at close to its rated speed), additional adapters continue to provide an incremental performance increase. Once the system reaches its performance limit, adding more adapters does not provide an increase in I/O throughput.

A number of factors can determine the performance limit. Bus speed, memory speed, adapter design, or processor speed can influence performance. Quite often, the system processor's speed may limit how many adapters of a given type the system can support while maintaining maximum performance. Once a

system uses 90 percent of its system processor, adding more adapters only provides a minor throughput increase.

Due to the wide variety of workloads, this section provides performance-limit guidelines only. The guidelines are based on I/O streaming of large reads or writes to a disk or network. They are not based on small I/Os, which are more transaction-rate limited. Small I/O workloads probably use more system processor capacity and result in fewer supported adapters for maximum performance.

The section bases these guidelines on the maximum number of processors supported for multi-processor systems.

If your system runs less than the maximum number of processors supported, then typically you must 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 your system uses disk and communication adapters concurrently, use a more conservative estimate of the number of supported adapters. If your configured system runs close to its performance limits, take extra care to ensure that your system type or configuration provides the desired performance. In these cases, you may need to contact your marketing support personnel for more detailed information.

Other restrictions

You must install some adapters in specific PCI slots in various systems. Physical size limits, I/O address considerations, thermal limitations, and other factors influence these specifications. This section lists slot placement information for PCI adapters in system units that have specific restrictions or guidelines. However, this book does not list all system-compatible PCI adapters.

7025 pSeries 620 Models 6F0 and 6F1

Figure B-1 on page 790 shows a system 7025 6F0 and 6F1 unit rear view with numbered slots




Figure B-1 System 7025 6F0 and 6F1 unit rear view with numbered slots

The 7025 Model 6F0 or 7025 Model 6F1 can accommodate up to 10 feature adapters based on the PCI bus. All slots are 64 bit with slots 6, 7, 11, and 12 supplying +5V and running at 33 MHz. The remaining slots are 3.3V slots capable of running at 66 MHz. The slots are numbered 3-12 (C03-C12), starting from the bottom of the machine.

Adapter cards that require +5V supply to operate must be plugged into slots 6, 7, 11, or 12. Adapter cards requiring +3.3V supply must be plugged into slots 3, 4, 5, 8, 9, or 10. Adapter cards that are universal (that is, they run on either voltage) can be plugged into any of the 10 slots. Some adapters must be placed in specific system unit slots to function correctly at optimum performance.

Use Table B-1 to identify adapter slot location options for the following adapters in your 7025 Model 6F0 or 7025 Model 6F1.

Table B-1 Placement guidelines 7025 6F0 or 6F1

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
2830	POWER GXT130P (Type 1-T)	6, 11, 7, 12	1	N

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	6, 11, 7, 12	1	N
2849	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	6, 11, 7, 12	1	N
2669	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6 (See notes 4 and 8)	Y
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6 (See notes 4 and 8)	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	9, 8, 5, 10, 4, 3	2	Y

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V	9, 8, 5, 10, 4, 3	2	Y
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12 (See note 3)	4	N
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	N
2708	ISDN Basic Rate PCI (Type 9-N)	6, 11, 7, 12	1	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6 (See note 13)	Y
6227	Gigabit Fibre Channel PCI (Type 4-S)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6 (See note 13)	N
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6235 is the Fast Write option for 6225 and 6230 FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM	9, 8, 3, 4, 10, 5, 6, 11, 7, 12 (See notes 1 and 2)	6	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6 (See notes 7 and 8)	Y

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	6, 11, 7, 12	4	N
4953	64-bit/66MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See note 8)	Y
4957	64-bit/66MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See note 8)	Y
2963	TURBOWAYS 155 PCI UTP ATM (Type 9-J)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See note 8)	Y
2988	TURBOWAYS 155 PCI MMF ATM (Type 9-F)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See note 8)	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See note 10)	Y
2742	SysKonnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See notes 5 and 8)	Y
2741	SysKonnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See notes 5 and 8)	Y

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
2743	SysKonnect SK-NET FDDI-UP SAS PCI (Type *)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	4 (See notes 5 and 8)	Y
2751	S/390 ESCON Channel, 32-bit, 5V, (Type 5-5)	6, 11, 7, 12	3	Y
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y
6205	PCI Dual Channel Ultra2 SCSI (Type 4-R)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	6	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	9	Y
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	6, 11, 7, 12	4	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10	Y
2968	10/100 Mbps Ethernet Tx PCI (Type 9-P)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10 (See notes 6 and 8)	Y

Feature code	Adapter	Slot usage	System maximum	Hot-pluggable
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12	10 (See notes 6 and 8)	Y
2985	Ethernet 10base2 PCI (Type 8-Y)	6, 11, 7, 12	4	Y
2987	Ethernet 10base5 PCI (Type 8-Z)	6, 11, 7, 12	4	Y
6311	IBM ARTIC960 RxF PCI, 32-bit, 3.3/5V (Type 6-G)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12 (See note 3)		
4958	Cryptographic Coprocessor (Type 6-H)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12		See note 11
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12		
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	9, 8, 3, 4, 10, 5, 6, 11, 7, 12		See note 12

Notes:

1. FC 6231: 128 MB DRAM Option Card for FC 6225 and FC 6230 is the 128 MB memory DIMM option for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter, and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max 1 per 6230).
2. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a RAID cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max 1 per 6230).

3. If multiple FC 6310s and FC 6311s are on the order, FC 2877 (Quad DTA, H.100, 4-Drop Cable) or FC 4353 (Cable, Internal H.100 BUS, 8-Position) must be used to connect FC 6310: IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E) and FC 6311: IBM ARTIC960 RxF PCI, 32-bit, 3.3/5V (Type 6-G).
4. For optimum system performance, a maximum of two FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U) or FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) adapters per system is recommended.
5. For optimum system performance, a maximum of 14 FC 2741: SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *), FC 2742: SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *), and FC 2743: SysKonnnect SK-NET FDDI-UP SAS PCI (Type *) adapters per system and a maximum of seven FC 2741, FC 2742, and FC 2743 adapters per I/O drawer is recommended.
6. For optimum system performance, there should be a maximum of eight FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters per system.
7. For optimum system performance using the FC 4961: 10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E) on this system, use only two of the four ports if the ports will all be running at media speed. Use a maximum of six ports per system (that is, three adapters running two ports each).
8. For optimum system performance, the combination of FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U), FC 2741: SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *), FC 2742: SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *), FC 2743: SysKonnnect SK-NET FDDI-UP SAS PCI (Type *), FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P), FC 4951: 10/100 4-Port Ethernet (Type 9-Z), FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J), FC 2988: TURBOWAYS 155 PCI MMF ATM (Type 9-F), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D), FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F), and FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) adapters cannot exceed seven adapters per system.
9. Any combination of the following adapters may be installed, but the combined total should not exceed four per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.

10. For optimum performance using FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B), limit adapters to the following:
 - MTU 1500: One per drawer, two per system
 - MTU 9180: Two per drawer, three per system
11. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
12. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
13. Use of one FC 6227: Gigabit Fibre Channel PCI (Type 4-S) or FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per bus is recommended. More than one of these adapters can be installed per bus, but system performance will not be increased by installing additional adapters to a bus.
14. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7026 pSeries 640 Model B80

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 B-2 shows a system 7026 B80 unit rear view with numbered slots.




Figure B-2 System 7026 B80 unit rear view with numbered slots

Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use Table B-2 to identify specific slot location options for the following adapters in your 7026 Model B80 system. 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 B-2 Placement guidelines 7026 B80

Feature code	Adapter	Usage	System maximum
2830	POWER GXT130P (Type 1-T)	5, 4, 3	1
4951	10/100 4-Port Ethernet (Type 9-Z)	4, 3	2
4958	Cryptographic Coprocessor (Type 6-H)	2, 1	2
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	2, 1	2
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	4, 3, 2, 1, 5	2 (See note 9)
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	4, 3, 2, 1, 5	2 (See notes 8 and 9)
6227	Gigabit Fibre Channel PCI (Type 4-S)	5, 4, 3, 2, 1	2 (See notes 7 and 9)
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	3, 4, 5	3
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 2, 3, 4	4
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 2, 3, 4	4
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	1, 2, 3, 4	2

Feature code	Adapter	Usage	System maximum
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	1, 2, 3, 4, 5	5 (See notes 5 and 3)
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	1, 2, 3, 4	4
6205	PCI Dual Channel Ultra2 SCSI (Type 4-R)	1, 2, 3, 4	4
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	5, 4, 3	2
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	1, 2, 3, 4	4
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 w/128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	4, 3, 2, 1 (See notes 5 and 6)	2
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	4, 3, 2, 1	4
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	4, 3, 2, 1	4
6311	IBM ARTIC960 RxF PCI, 32-bit, 3.3/5V (Type 6-G)	4, 3, 2, 1	3
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	4, 3, 2, 1	4 (See notes 1 and 4)
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	5, 4, 3, 2, 1	5 (See note 2)
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	5, 4, 3, 2, 1	1
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	5, 4, 3, 2, 1	1
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	5, 4, 3, 2, 1	5

Feature code	Adapter	Usage	System maximum
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	5, 4, 3, 2, 1	5
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	5, 4, 3, 2, 1	5 (See note 1)
4953	64-bit/66MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	5, 4, 3, 2, 1	5
4957	64-bit/66MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	5, 4, 3, 2, 1	5
2963	TURBOWAYS 155 PCI UTP ATM (Type 9-J)	5, 4, 3, 2, 1	5
2988	TURBOWAYS 155 PCI MMF ATM (Type 9-F)	5, 4, 3, 2, 1	5
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	5, 4, 3, 2, 1	5
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	5, 4, 3, 2, 1	5
2743	SysKonnnect SK-NET FDDI-UP SAS PCI (Type *)	5, 4, 3, 2, 1	5
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	5, 4, 3, 2, 1	5
2968	10/100 Mbps Ethernet Tx PCI (Type 9-P)	5, 4, 3, 2, 1	5 (See note 2)
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	5, 4, 3, 2, 1	5
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	5, 4, 3, 2, 1	5 (See note 1 and 3)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	5, 4, 3, 2, 1	4

Notes:

1. For optimum system performance, a combination of the following adapters should not exceed a maximum of two adapters per system:
 - FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)

- FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
- FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)

Install one in slots 1 or 2 and the other in slots 3 or 4.

2. This system supports up to five FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters per system, making the total six if you use the Integrated Ethernet port.
3. For optimum system performance, install a maximum of two FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) or two FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U) adapters per system (or one of each), when configured to run at 1000 Mbps. If the adapter is configured for 100 Mbps mode, use up to four of these adapters (plus the two integrated 100 Mbps ports) for optimum performance.
4. For optimum system performance, install a maximum of two FC 4951: 10/100 4-Port Ethernet (Type 9-Z) adapters in 100 Mbps mode per system.
5. FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) can be selected with or without the cache upgrade (FC 6235: Fast Write Cache Option for FC 6225 and FC 6230).
6. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (maximum one per 6230).
7. If two FC 6227: Gigabit Fibre Channel PCI (Type 4-S) adapters are to be used, the first must be placed in slots 1 or 2, and the second must be placed in slots 3, 4, or 5.
8. FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) will operate in 32-bit slots, but at a decrease in performance.
9. Use of one FC 6227: Gigabit Fibre Channel PCI (Type 4-S), FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W), or FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) per bus is recommended. More than one of these adapters can be installed per bus, but system performance will not be increased by installing additional adapters to a bus.
10. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7028 pSeries 610 Models 6C1 and 6E1

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 B-3 shows a system 7028 6C1 and 6E1 unit rear view with numbered slots.




Figure B-3 System 7028 6C1 and 6E1 unit rear view with numbered slots

Some adapters must be placed in specific system unit slots to function correctly at optimum performance.

Use Table B-3 on page 804 to determine where to install an adapter in your system unit.

Many of the following notes refer to optimizing system performance. Read “System performance” on page 787 for more performance-related information.

Use the rear-view diagram above to identify slot locations described in Table B-3 on page 804.

Table B-3 Placement guidelines 7028 6C1 and 6E1

Feature code	Adapter	Slot usage	Maximum per drawer/system
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	1, 2	2 (See note 9)
2830	POWER GXT130P (Type 1-T)	5, 4, 3	2
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	1, 3, 2, 4, 5	2 to 4 (See notes 8 and 10)
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	1, 3, 2, 4, 5	2 to 4 (See notes 8 and 10)
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 2	2
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 2	2
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	1, 2, 3, 4, 5	5 (See note 4)
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	1, 2, 3, 4, 5	5 (See note 4)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5

Feature code	Adapter	Slot usage	Maximum per drawer/system
2849	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 2849)	1, 2, 3, 4, 5	4
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1, 2, 3, 4, 5	2
2751	S/390 ESCON Channel, 32-bit, 5V, (Type 5-5)	3, 4, 5, 1	2
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	5, 4, 3, 2, 1	2
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	1, 2, 3, 4, 5	5
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	1, 2, 3, 4, 5	5
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	1, 2, 3, 4, 5	5
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	1, 2, 3, 4, 5	5

Feature code	Adapter	Slot usage	Maximum per drawer/system
6230	Advanced SerialRAID Plus 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	1, 2, 3, 4, 5	2 (See notes 1 and 2)
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	5, 4, 3	2
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	5, 4, 3, 2, 1	4
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	1, 2, 3, 4, 5	4
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	1, 2, 3, 4, 5	4
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	1, 2, 3, 4, 5	2
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	1, 2, 3, 4, 5	2
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	5, 4, 3, 2, 1	5 (See notes 5 and 7)

Feature code	Adapter	Slot usage	Maximum per drawer/system
2943	8-Port Asynchronous EIA-232E/RS-422 A PCI, 32-bit, 3.3/5V (Type 3-B)	5, 4, 3, 2, 1	5
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	5, 4, 3, 2, 1	5
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	5, 4, 3, 2, 1	4
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	5, 4, 3, 2, 1	5
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	5, 4, 3, 2, 1	5
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	5, 4, 3, 2, 1	5 (See note 5)
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	5, 4, 3, 2, 1	5 (See note 5)
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 2, 3, 4, 5	5 (See note 4)

Notes:

1. FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) or FC 6230+6231: 4-Port SSA 40w/128 MB DIMM can be selected with or without the cache upgrade (FC 6235: Fast Write Cache Option for FC 6225 and FC 6230).

2. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a cache upgrade for the FC 6230 + 6231: 4-Port SSA 40 with 128 MB DIMM adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) or FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM (max 1 per FC 6230/FC 6230+6231).
3. FC 6227: Gigabit Fibre Channel PCI (Type 4-S) is restricted to only one adapter in slot 1 or 2 and/or one adapter in slot 3, 4 or 5.
4. For optimum system performance, a combination of the following adapters should not exceed a maximum of one adapter per system: FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U), FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A), and FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B). Use 64-bit slots (1, 2, or 3) if available.
5. For optimum system performance, install up to five FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C) or FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) adapters. Install up to three of these adapters if you are in MTU 1500 mode. Use 64-bit slots (1, 2, or 3) if available.
6. For optimum system performance, install up to three FC 4962: Ethernet/LAN Encryption10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters in this system.
7. For optimum system performance, a maximum of one FC 4961: 10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E) adapter per system (with a maximum of three ports) is recommended.
8. FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) will operate in 32-bit slots, but at a decrease in performance. If RPQ# 8A1209 is in the system, the number of FC 6228 adapters in the system can be increased to four, but system performance will decrease.
9. Any combination of the following adapters may be installed, but the combined total should not exceed two per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
10. One FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
11. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7028 pSeries 630 Models 6C4 and 6E4 (4-slot PCI riser)

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. Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system.

In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error. To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Figure B-4 on page 810 shows a system 7028 6C4 and 6E4 (4-Slot PCI riser) unit rear view with numbered slots.




Figure B-4 7028 6C4 and 6E4 rear view (4-slot PCI riser) with numbered slots

Table B-4 provides the Slot Location Reference 7028 6C4 and 6E4 (4-Slot PCI riser card).

Table B-4 Slot Location Reference 7028 6C4 and 6E4 (4-Slot PCI riser card)

Slot location	PHB	Planar	Loc. Code	Characteristics
1	1	1	Ux.y-P2-I1	64-bit 3.3V, 133 MHz
2	1	1	Ux.y-P2-I2	64-bit 3.3V, 133 MHz
3	0	1	Ux.y-P2-I3	64-bit 3.3V, 133 MHz
4	0	1	Ux.y-P2-I4	64-bit 3.3V, 133 MHz

In Table B-4 Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Table B-5 on page 811 is to identify specific slot location options for the adapters in the 4-Slot PCI riser of your 7028 Models 6C4 and 6E4 (pSeries 630) system.

Table B-5 Placement guidelines 7028 6C4 and 6E4 (4-slot PCI riser card)

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
8397	SP System Attachment, 64-bit, 3.3/5V, 2-slot	3 (See note 4)	1/1	1	Y	Y
2842	POWER GXT4500P Graphics Adapter (Type 1-Y)	4, 3	1/1 (See note 6)	1	Y	N
2843	POWER GXT6500P Graphics Adapter (Type 1-Z)	4, 3	1/1 (See note 6)	1	Y	N
8244	Audio PCI Adapter for Workstations (Type 8244)	1, 2 (See note 8)	1/1	1	N	N
6312	Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312)	4, 3, 2, 1	6/16	4	Y	Y
6310	IBM ARTIC960 Red Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	4, 3, 2, 1	4/16	4	N	N
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	4, 3, 2, 1	4/4	4	Y	Y (See note 2)
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	3, 2, 1	3/17	17	Y	Y
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	4, 3, 2, 1	4/18	18	Y	Y (See note 2)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	4, 3, 2, 1	4/18	18	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	4, 3, 2, 1	4/18	18	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	4, 3, 2, 1	4/18	18	Y	Y

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	4, 3, 2, 1	4/18	18	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	4, 3, 2, 1	4/4	4	Y	Y
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	4, 3, 2, 1	4/4	4	Y	Y
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	4, 3, 2, 1	4/4	4	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	4, 3, 2, 1	4/18 (See note 3)	18	Y	Y
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	4, 3, 2, 1	4/18	18	Y	Y
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	4, 3, 2, 1	4/18	18	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	4, 3, 2, 1	4/18	18	N/A	N/A
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	4, 3, 2, 1	4/18	18	Y	Y
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	4, 3, 2, 1	4/18	18	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	4, 3, 2, 1	4/18	18	Y	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	4, 3, 2, 1	4/18	18	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	4, 3, 2, 1	4/4	1	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	4, 3, 2, 1	4/4	1	Y	N

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	4, 3, 2, 1	4/10	10	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	4, 3, 2, 1	4/18	18	Y	Y
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	4, 3, 2, 1	2/2 (See note 7)	1	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	4, 3, 2, 1	2/2 (See note 7)	1	Y	Y
4953	64-bit/66MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	4, 3, 2, 1	4/10	10	Y	Y
4957	64-bit/66MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	4, 3, 2, 1	4/10	10	Y	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	4, 3, 2, 1	4/14	14 (See note 5)	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	4, 3, 2, 1	4/18	18	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	4, 3, 2, 1	4/18	18	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	4, 3, 2, 1	4/4	1	Y	N
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	4, 3, 2, 1	4/6	6	N	Y
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	4, 3, 2, 1	4/6	6	N	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	4, 3, 2, 1	4/18	18	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	4, 3, 2, 1	4/18	18	Y	Y

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	4, 3, 2, 1	4/18	5 (See note 5)	Y	N

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
4. Use of FC 8397: SP System Attachment, 64-bit, 3.3/5V, 2-slot is for the 7028 Model 6C4 only. If the adapter is installed in slot 3, no adapter can be installed in slot 4, and the divider between slot 3 and slot 4 must be removed.
5. The total of any combination of FC 2947: IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R) and FC 2962: 2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V) must not exceed 14 per LPAR.
6. The total of FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y) and FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z) must not exceed one per riser and one per system. If one 3D graphics adapter is installed in either slot 3 or 4, no adapter can be installed in other slot of this pair.
7. The total of any combination of FC 2732: Serial HIPPI, Short-Wave Optics (Type 9-W) and FC 2733: Serial HIPPI, Long-Wave Optics (Type 9-W) must not exceed two per system.
8. Use of FC 8244: Audio PCI Adapter for Workstations (Type 8244) is for the 7028 Model 6E4 only.
9. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7028 pSeries 630 Models 6C4 and 6E4 (6-Slot PCI Riser)

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. Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Figure B-5 shows a system 7028 6C4 and 6E4 (6-Slot PCI riser) unit rear view with numbered slots.




Figure B-5 7028 6C4 and 6E4 rear view (6-slot PCI riser) with numbered slots

All slots in the 6-Slot PCI Riser are capable of using full-length PCI-X adapters.

Table B-6 show the Slot Location Reference 7028 6C4 and 6E4 (6-Slot PCI riser card).

Table B-6 Slot Location Reference 7028 6C4 and 6E4 (6-Slot PCI riser card)

Slot Loc.	PHB	Planar	Loc. Code	Characteristics
1	1	1	Ux.y-P2-I1	64-bit 3.3V, 66 MHz
2	1	1	Ux.y-P2-I2	64-bit 3.3V, 66 MHz
3	2	1	Ux.y-P2-I3	64-bit 3.3V, 133 MHz
4	2	1	Ux.y-P2-I4	64-bit 3.3V, 133 MHz

Slot Loc.	PHB	Planar	Loc. Code	Characteristics
5	0	1	Ux.y-P2-I5	64-bit 3.3V, 133 MHz
6	0	1	Ux.y-P2-I6	64-bit 3.3V, 133 MHz

Note: In Table B-6, Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Table B-7 is to identify specific slot location options for the adapters in the 6-Slot PCI riser of your 7028 Models 6C4 and 6E4 (pSeries 630) system.

Table B-7 Placement guidelines 7028 6C4 and 6E4 (6-Slot PCI riser card)

Feature code	Adapter	Slot usage	Max per riser/ system	Max per LPAR	EEH	Hot-plug
8397	SP System Attachment, 64-bit, 3.3/5V, 2-slot	5 (See note 3)	1/1	1	Y	Y
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot	3 (See note 7)	1/2	1	Y	N
2842	POWER GXT4500P Graphics Adapter (Type 1-Y)	3, 4	1/1 (See note 4)	1	N	N
2843	POWER GXT6500P Graphics Adapter (Type 1-Z)	3, 4	1/1 (See note 4)	1	N	N
8244	Audio PCI Adapter for Workstations (Type 8244)	1,2 (See note 6)	1	1	N	N
6312	Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312)	3, 4, 5, 6, 1, 2	6/16	4	Y	Y
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	3, 4, 5, 6, 1, 2	6/16	4	N	N
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	3, 4, 5, 6, 1, 2	6/16	4	Y	Y (See note 1)

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y (See note 2)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	5, 6, 4, 3, 1, 2	6/20	20	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	5, 6, 4, 3, 1, 2	6/20	20	Y	Y
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	3, 4, 5, 6, 1, 2	6/6	6	Y	Y
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	3, 4, 5, 6, 1, 2	4/4	4	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	3, 4, 5, 6, 1, 2	4/18 (See note 8)	18	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	3, 4, 5, 6, 1, 2	4/4	4	Y	Y
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	3, 4, 5, 6, 1, 2	4/18	18	Y	Y
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	3, 4, 5, 6, 1, 2	4/18	18	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	3, 4, 5, 6, 1, 2	4/18	18	N/A	N/A
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	3, 4, 5, 1, 2	3/11	11	Y	Y

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	3, 4, 5, 6, 1, 2	6/6	6	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	3, 4, 5, 6, 1, 2	6/20	20	N	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
2849	FC 2849: POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	3, 4, 5, 6, 1, 2	4/4	1	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	3, 4, 5, 6, 1, 2	4/4	1	Y	N
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	3, 4, 5, 6, 1, 2	5/19	19	Y	Y
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	3, 4, 5, 6, 1, 2	5/19	19	Y	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	3, 4, 5, 6, 1, 2	6/20	16	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	3, 4, 5, 6, 1, 2	2/2 (See note 5)	2	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	3, 4, 5, 6, 1, 2	2/2 (See note 5)	2	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	3, 4, 5, 6, 1, 2	6/20	16	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	3, 4, 5, 6, 1, 2	6/20	16	Y	Y

Feature code	Adapter	Slot usage	Max per riser/system	Max per LPAR	EEH	Hot-plug
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	3, 4, 5, 6, 1, 2	4/4	1	Y	N
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	3, 4, 5, 6, 1, 2	4/6	6	N	Y
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	3, 4, 5, 6, 1, 2	4/6	6	N	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	3, 4, 5, 6, 1, 2	6/20	20	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. Use of FC 8397: SP System Attachment, 64-bit, 3.3/5V, 2-slot is for the 7028 Model 6C4 only. This adapter must be installed in slot 5, and no other adapter can be installed in slot 6. Remove the divider between slots 5 and 6 before installing this adapter.
4. Use of FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y) and FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z) is for the 7028 Model 6E4 only. The total of FC 2842: POWER GXT4500P Graphics Adapter (Type

- 1-Y) and FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z) must not exceed one per riser and one per system.
5. The total of any combination of FC 2732: Serial HIPPI, Short-Wave Optics (Type 9-W) and FC 2733: Serial HIPPI, Long-Wave Optics (Type 9-W) must not exceed not exceed two per system.
 6. Use of FC 8244: Audio PCI Adapter for Workstations (Type 8244) is for the 7028 Model 6E4 only.
 7. Use of FC 8398: SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot is for 7028 Model 6C4 only.
 8. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
 9. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7029 pSeries 615 Models 6C3 and 6E3

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. Some adapters must be placed in specific system unit slots to function correctly at optimum performance.

Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB

from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Figure B-6 shows a system 7029 6C3 and 6E3 unit rear view with numbered slots.




Figure B-6 System 7029 6C3 and 6E3 unit rear view with numbered slots

All slots in the system are capable of using PCI-X adapters. Slots 1, 4, 5, and 6 support full-length PCI-X adapters. Slots 2 and 3 support half-length PCI-X adapters.

Table B-8 show the Slot Location Reference 7029 6C3 and 6E3.

Table B-8 Slot Location Reference 7029 6C3 and 6E3

Slot Loc.	PHB	Planar	Loc. Code	Characteristics
1	1	1	Ux.y-P1-I1	64-bit 3.3V, 66 MHz
2	1	1	Ux.y-P1-I2	32-bit 3.3V, 66 MHz

Slot Loc.	PHB	Planar	Loc. Code	Characteristics
3	1	1	Ux.y-P1-I3	32-bit 3.3V, 133 MHz
4	1	1	Ux.y-P1-I4	64-bit 3.3V, 133 MHz
5	1	1	Ux.y-P1-I5	64-bit 3.3V, 133 MHz
6	1	1	Ux.y-P1-I6	64-bit 3.3V, 133 MHz

Note: In Table B-8, Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Table B-9 to identify specific slot location options for the adapters in your 7029 Models 6C3 and 6E3 system.

Table B-9 Placement guidelines 7029 6C3 and 6E3

Feature code	Adapter	Slot usage	Max per system	EEH	Hot-plug
5709	Dual Channel SCSI RAID Enablement Card (Type 5709)	N/A	1	N/A	N
6312	Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312)	4, 5, 6 (See note 3)	3	Y	Y
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	4, 5, 6 (See note 3)	3	N	N
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	4, 5, 6	3	Y	Y
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	4, 5, 6	3	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	4, 5, 6	3	N/A	N/A
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	4, 5, 6	3	Y	Y
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	4, 5, 6	3	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	4, 5, 6	3	Y	Y
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	4, 5, 6	3	Y	Y

Feature code	Adapter	Slot usage	Max per system	EEH	Hot-plug
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 4, 5, 6	4	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 4, 5, 6	4	Y	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 4, 5, 6	4	Y	Y
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	1, 4, 2, 3, 5, 6	6	Y	Y (See note 1)
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 4, 5, 6	4	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 4, 5, 6	4	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	1, 4, 5, 6	4	Y	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	1, 4, 2, 3, 5, 6	6	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 4, 2, 3, 5, 6	4	Y	N
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	4, 5, 6	3	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	4, 5, 6	3	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	4, 5, 6	2	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 4, 2, 3, 5, 6	6	Y	Y
8244	Audio PCI Adapter for Workstations (Type 8244)	1, 4, 2, 3, 5, 6 (See note 4)	1	N	N
4953	64-bit/66MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	1, 4, 5, 6	4	Y	Y

Feature code	Adapter	Slot usage	Max per system	EEH	Hot-plug
4957	64-bit/66MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	1, 4, 5, 6	4	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 4, 2, 3, 5, 6	6	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	1, 4, 2, 3, 5, 6 (See note 2)	1	Y	N
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 4, 2, 3, 5, 6	6	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 4, 2, 3, 5, 6	6	Y	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 4, 2, 3, 5, 6	6	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
2. The FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V is only available for Linux operating systems.
3. Digital Trunk adapters have an internal cable and must be in contiguous slots.
4. Use of FC 8244: Audio PCI Adapter for Workstations (Type 8244) is for 7029 Model 6E3 only.
5. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7040 pSeries 670 Model 671 drawers 61D

Adapter cards for the 7040 Model 61D (pSeries 670) plug into PCI adapter slots in the left or right side of the I/O subsystem (some system configurations contain only the left side PCI adapter slots). Each drawer is capable of handling up to 20 PCI adapters. Slots 1 through 7 and 11 through 17 have 3.3V signaling and can support 33 MHz or 66 MHz operation depending on the adapter's capability. 5V only adapters cannot be used in these slots. Slots 8, 9, 10 and 18, 19, 20 have

5V signaling and can only run at 33 MHz. 5V only adapters must use these slots. All slots provide both 3.3V and 5V power for the adapters.

Figure B-7 shows a system 7040 671 drawer 61D unit rear view with numbered slots.




Figure B-7 System 7040 671 drawer 61D unit rear view with numbered slots

Table B-10 shows the Slot Location Reference 7040 pSeries 670 drawer 61D.

Table B-10 Slot Location Reference 7040 pSeries 670 drawer 61D

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
1	1	1	Ux.y-P1-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
4	1	1	Ux.y-P1-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
5	2	1	Ux.y-P1-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
6	2	1	Ux.y-P1-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
7	2	1	Ux.y-P1-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	1	Ux.y-P1-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
8	3	1	Ux.y-P1-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
9	3	1	Ux.y-P1-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
10	3	1	Ux.y-P1-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
See note 1	3	1	Ux.y-P1-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
11	1	2	Ux.y-P2-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
12	1	2	Ux.y-P2-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
13	1	2	Ux.y-P2-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
14	1	2	Ux.y-P2-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
15	2	2	Ux.y-P2-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
16	2	2	Ux.y-P2-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
17	2	2	Ux.y-P2-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	2	Ux.y-P2-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
18	3	2	Ux.y-P2-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
19	3	2	Ux.y-P2-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
20	3	2	Ux.y-P2-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
See note 1	3	2	Ux.y-P2-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz

Notes:

1. Integrated SCSI adapter.
2. Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Figure B-7 on page 825 to identify the slot locations described in Table B-11 on page 827. The first adapter from an adapter group is placed in the first slot (per the table) in the drawer. After the first slot selection has been used, use the second slot in the list. If a designated slot is filled, move to the next available slot.

Note: The default configuration of a 7040 pSeries 670 Model 61D's I/O Subsystem contains only one Planar in the drawer. If your system contains only Planar 1, then ignore the slot restrictions listed in Table B-11 for Planar 2 (P2-I).

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Use Table B-11 to identify specific slot locations for the following adapters in your 7040 pSeries 670 Model 61D. The slots listed refer to the HMC (Hardware Management Console) location codes detailed in Figure B-7 on page 825.

Table B-11 Placement guidelines 7040 671 drawer 61D

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	EEH	Hot-Plug
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot	3, 13, 5, 15 (See note 19)	3, 5 (See note 19)	0/2	4	2	0/8	Y	N
8397	SP System Attachment, 64-bit, 3.3/5V, 2-slot	3+4, 13+14, 5+6, 15+16 (See note 20)	3+4, 13+14, 5+6, 15+16 (See note 20)	2/0	4	2	8/0	Y	Y
8396	SP System Attachment, 32-bit, 5V, 2-slot (Type 6-F)	8+9, 18+19 (See note 20)	18+19 (See note 20)	1/0	2	1	4/0	N	N

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	EEH	Hot-Plug
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	10, 20 (See note 20)	20 (See note 20)	1/0	2	2	2/0 (See notes 4, 10, and 11)	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	Two FC 6563: 10, 20, 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 10, 20, 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 4, 14	10, 20, 1, 5, 8, 2, 6, 9, 3, 7, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30 (See notes 4 and 11)	Y	Y
5710	PCI- X Dual Channel Ultra320 SCSI Blind Swap Adapter, 32/64-bit, 3.3V (Type 5710)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 10, 20 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	20	30	Y	Y
5711	PCI- X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter, 32/64-bit, 3.3V (Type 5711)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 10, 20 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	20	30	Y	Y
2751	S/390 ESCON Channel, 32-bit, 5V, (Type 5-5)	9, 19, 10, 20 (See note 20)	19, 20 (See note 20)	2/0	4	8	8/0	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	50	30/50	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	EEH	Hot-Plug
6228	FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) (See note 23)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	50	30/50 (See note 15)	Y	Y
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30 (See note 3)	Y	Y
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30 (See note 3)	Y	Y
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	Two FC 6563: 1, 11, 3, 13, 4, 14, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14 (See notes 1, 2, 12, and 13)	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19 (See notes 1, 2, 12, and 13)	10	20	30	30	Y	Y
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30/60	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30/60	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	20/40	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	20/40	Y	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30 (See note 5)	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	(5,6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 6)	(5,6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 6)	2	4	1	8	Y	N

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	EEH	Hot-Plug
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	(5, 6, 7, 15, 16, 17); (1, 2, 3, 4, 11, 12, 13, 14); or (8, 9, 10, 18, 19, 20) (See note 6)	(5, 6, 7, 15, 16, 17); (1, 2, 3, 4, 11, 12, 13, 14); or (8, 9, 10, 18, 19, 20) (See note 6)	2	4	1	8	Y	N
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	2	4	2	8 (See note 7)	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	2	4	2	8 (See note 7)	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20 (See note 13)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20 (See note 13)	10	20	60	60	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	5/10	10/20	20	20/40	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	32	32 (See note 8)	Y	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4	8	16	16 (See notes 8 and 14)	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40 (See note 5)	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40 (See note 5)	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40	Y	Y
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	8	8	8	8 (See note 9)	N	Y
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	8	8	8	8 (See note 9)	N	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	Two FC 6563: 10, 20, 9, 19, 8, 18 Two FC 6571: 10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 5, 4, 3, 2, 1	3/6	6	18	18 (See note 21)	Y	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	10, 20, 9, 19, 7, 17, 6, 16, 4, 14, 3, 13, 2, 12	10, 20, 9, 19, 7, 17, 6, 16, 4, 14, 3, 13, 2, 12	10	16	14	16	Y	Y
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4	8	24	24	Y	Y (See note 17)
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-l)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4 (See note 18)	8 (See note 18)	4 (See note 18)	24 (See note 18)	Y	Y (See note 16)
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	5	10	20	20 (See note 10)	Y	Y

Notes:

1. FC 6231: 128 MB DRAM Option Card for FC 6225 and FC 6230 is a 128 MB memory DIMM option for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter, and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max of one per 6230).
2. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a RAID cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max of one per 6230).
3. A combined maximum of 30 FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U) and FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) adapters is allowed per system.
4. FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) with FC 2118 or FC 6206: Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K) is required with FC 4609 to drive the media subsystem with up to two media devices. FC 6206 has first priority. If FC 6203 is used, then FC 2118 will be placed with the drawer.
5. A combined maximum of 40 FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), and FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B) adapters is allowed per system.
6. FC 2849: POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849), FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X), and FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V must be placed on the same PHB. All other adapters on that PHB must be on the same LPAR. In the table, the slots of each PHB are divided into parentheses for reference.
7. A combined maximum of eight FC 2732: Serial HIPPI, Short-Wave Optics (Type 9-W) and FC 2733: Serial HIPPI, Long-Wave Optics (Type 9-W) adapters is allowed per system.
8. A combined maximum of 32 FC 2943: 8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B) and FC 2944: 128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C) adapters is allowed per system.
9. A combined maximum of eight FC 2741: SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *), and FC 2742: SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *) adapters is allowed per system.
10. A maximum of 20 FC 6204: PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U) and FC 6206: Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K) adapters is allowed per system.

11. FC 6206: Single-Ended Ultra SCSI (Type 4-K) may be utilized to connect media devices located in the media subsystem. If no FC 6206 is present in the system, FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) may be used to connect the media devices in the media subsystem.
12. A maximum of three SSA adapters are allowed in PHB1 of each planar (slots P1-I1, P2-I1, P1-I3, P2-I3, P1-I4, and P2-I4).
13. If three FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) adapters are populated in PHB1, then only FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) may be populated in the remaining slot of that PHB. It must be placed in slots P1-I2 or P2-I2. This allows for six high-performance adapters if the first three in each IO planar are FC 6230.
14. When installing FC 2943: 8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B) in slots P1-I1, P2-I1, P1-I10, and P2-I10, the metal tab on the adapter should be pushed down to assist in plugging.
15. When installing FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) in slot 1, the metal tab on the adapter should be pushed down to assist in plugging.
16. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
17. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
18. Any combination of the following adapters may be installed, but the combined total should not exceed four per planar or LPAR, eight per drawer, or 24 per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
19. The slot usage applies to only FC 6571: I/O Drawer PCI-X Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports.
20. The slot usage applies to only FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports.
21. If only one FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports is installed in the system, FC 2962: 2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V) is limited to a maximum of three per LPAR.
22. High performance adapters that are installed in FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports have the following limitations:
 - Two per first PHB on each planar
 - two per second PHB on each planar
 - One per third PHB on each planar

23. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
24. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7040 pSeries 690 Model 681 drawer 61D

Adapter cards for the 7040 pSeries 690 Model 61D plug into PCI adapter slots in the left or right side of the I/O subsystem (some system configurations contain only the left side PCI adapter slots). Each drawer is capable of handling up to 20 PCI adapters. Slots 1 through 7 and 11 through 17 have 3.3V signaling and can support 33 MHz or 66 MHz operation depending on the adapter's capability. 5V only adapters cannot be used in these slots. Slots 8, 9, and 10 and 18, 19, and 20 have 5V signaling and can only run at 33 MHz. 5V only adapters must use these slots. All slots provide both 3.3V and 5V power for the adapters.

Figure B-8 shows a system 7040 681 drawer 61D unit rear view with numbered slots.




Figure B-8 System 7040 681 drawer 61D unit rear view with numbered slots

Table B-12 on page 836 show the Slot Location Reference for the 7040 pSeries 690 drawer 61D.

Table B-12 Slot Location Reference for the 7040 681 drawer 61D

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
1	1	1	Ux.y-P1-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
4	1	1	Ux.y-P1-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
5	2	1	Ux.y-P1-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
6	2	1	Ux.y-P1-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
7	2	1	Ux.y-P1-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	1	Ux.y-P1-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
8	3	1	Ux.y-P1-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
9	3	1	Ux.y-P1-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
10	3	1	Ux.y-P1-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
See note 1	3	1	Ux.y-P1-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
11	1	2	Ux.y-P2-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
12	1	2	Ux.y-P2-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
13	1	2	Ux.y-P2-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
14	1	2	Ux.y-P2-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
15	2	2	Ux.y-P2-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
16	2	2	Ux.y-P2-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
17	2	2	Ux.y-P2-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	2	Ux.y-P2-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
18	3	2	Ux.y-P2-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
19	3	2	Ux.y-P2-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
20	3	2	Ux.y-P2-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
See note 1	3	2	Ux.y-P2-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz

Notes:

1. Integrated SCSI adapter.
2. Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Figure B-8 on page 835 to identify slot locations described in Table B-13 on page 838. The first adapter from an adapter group is placed in the first slot (per the table) in the drawer. After the first slot selection has been used, use the second slot in the list. If a designated slot is filled, move to the next available slot.

Note: The default configuration of a 7040 pSeries 670 Model 61D's I/O Subsystem contains only one Planar in the drawer. If your system contains only Planar 1, then ignore the slot restrictions listed in Table B-13 for Planar 2 (P2-I).

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Use Table B-13 to identify specific slot locations for the following adapters in your 7040 pSeries 690 Model 61D. The slots listed refer to the HMC (Hardware Management Console) location codes detailed in the preceding illustration.

Table B-13 Placement guidelines 7040 pSeries 690 drawer 61D

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot	3, 13, 5, 15 (See note 19)	3, 5 (See note 19)	0/2	4	2	0/32	Y	N
8397	SP System Attachment, 64-bit, 3.3/5V, 2-slot	3+4, 13+14, 5+6, 15+16	3+4, 13+14, 5+6, 15+16	2/0	4	2	32/0	Y	Y
8396	SP System Attachment, 32-bit, 5V, 2-slot (Type 6-F)	8+9, 18+19 (See note 20)	18+19 (See note 20)	1/0	2	1	8/0	N	N
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	10, 20 (See note 20)	20 (See note 20)	1/0	2	2	2/0 (See notes 4,10,and 11)	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	Two FC 6563: 10, 20, 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 10, 20, 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 4, 14	10, 20, 1, 5, 8, 2, 6, 9, 3, 7, 4, 11, 13, 15, 17, 19	5/10	10/20	30	30 (See notes 4 and 11)	Y	Y
5710	PCI- X Dual Channel Ultra320 SCSI Blind Swap Adapter, 32/64-bit, 3.3V (Type 5710)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 10, 20 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	40	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
5711	PCI- X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter, 32/64-bit, 3.3V (Type 5711)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 10, 20 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	40	Y	Y
2751	S/390 ESCON Channel, 32-bit, 5V, (Type 5-5)	9, 19, 10, 20 (See note 20)	19, 20 (See note 20)	2/0	4	4	8/0	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80/140	Y	Y
6228	FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) (See note 23)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80 (See note 15)	Y	Y
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80 (See note 3)	Y	Y
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80 (See note 3)	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	Two FC 6563: 1, 11, 3, 13, 4, 14, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14 (See notes 1, 2, 12, and 13)	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19 (See notes 1, 2, 12, and 13)	10	20	64	64	Y	Y
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80/140	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	80	80/140	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	40/80	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	Two FC 6563: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 4, 14 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	40/80	Y	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	Two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 9, 19 Two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	1, 5, 8, 2, 6, 9, 3, 7, 10, 4, 11, 13, 15, 17, 19	5/10	10/20	40	40 (See note 5)	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	(5,6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 6)	(5,6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 6)	2	4	1	8	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	(5, 6, 7, 15, 16, 17); (1, 2, 3, 4, 11, 12, 13, 14); or (8, 9, 10, 18, 19, 20) (See note 6)	(5, 6, 7, 15, 16, 17); (1, 2, 3, 4, 11, 12, 13, 14); or (8, 9, 10, 18, 19, 20) (See note 6)	2	4	1	8	Y	N
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	2	4	2	8 (See note 7)	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	2	4	2	8 (See note 7)	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20 (See note 13)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20 (See note 13)	10	20	80	80/140	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 or one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	5/10	10/20	20	20/40	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	32	32 (See note 8)	Y	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4	8	16	16 (See notes 8 and 14)	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40 (See note 5)	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40 (See note 5)	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10	20	40	40	Y	Y
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	8	8	8	8 (See note 9)	N	Y
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	8	8	8	8 (See note 9)	N	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	Two FC 6563: 10, 20, 9, 19, 8, 18 Two FC 6571: 10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 5, 4, 3, 2, 1	3/6	6	20	20	Y	Y

Feature code	Adapter	Slot Usage when there are two FC 6563 or two FC 6571 Planars	Slot Usage when there is one FC 6563 and one FC 6571 Planar	Max. per FC 6563 or FC 6571 Planar	Max. per Drawer when there are two FC 6563 or two FC 6571 Planars	Max. per LPAR	Max. per System when there are two FC 6563 or two FC 6571 Planars	E E H	Hot-Plug
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	10, 20, 9, 19, 7, 17, 6, 16, 4, 14, 3, 13, 2, 12	10, 20, 9, 19, 7, 17, 6, 16, 4, 14, 3, 13, 2, 12	10	16	14	16	Y	Y
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4	8	32	32	Y	Y (See note 17)
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	4 (See note 18)	8 S(see note 18)	4 (See note 18)	32 (See note 18)	Y	Y (See note 16)
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	5	10	20	20 (See note 10)	Y	Y

Notes:

1. FC 6231: 128 MB DRAM Option Card for FC 6225 and FC 6230 is a 128 MB memory DIMM option for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max of one per 6230).
2. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a RAID cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (max of one per 6230).
3. A combined maximum of 30 FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U) and FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) adapters is allowed per system.
4. FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) with FC 2118 or FC 6206: Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K) is required with FC 4609 to drive the media subsystem with up to two media devices. FC 6206 has first priority. If FC 6203 is used, then FC 2118 will be placed with the drawer.

5. A combined maximum of 40 FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), and FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B) adapters is allowed per system.
6. FC 2849: POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849), FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) and FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V must be placed on the same PHB. All other adapters on that PHB must be on the same LPAR. In the table, the slots of each PHB are divided into parentheses for reference.
7. A combined maximum of eight FC 2732: Serial HIPPI, Short-Wave Optics (Type 9-W) and FC 2733: Serial HIPPI, Long-Wave Optics (Type 9-W) adapters is allowed per system.
8. A combined maximum of 32 FC 2943: 8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B) and FC 2944: 128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C) adapters is allowed per system.
9. A combined maximum of eight FC 2741: SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *), and FC 2742: SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *) adapters is allowed per system.
10. A maximum of 20 FC 6204: PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U) and FC 6206: Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K) adapters is allowed per system.
11. FC 6206: Single-Ended Ultra SCSI (Type 4-K) may be utilized to connect media devices located in the media subsystem. If no FC 6206 is present in the system, FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) may be used to connect the media devices in the media subsystem.
12. A maximum of three SSA adapters are allowed in PHB1 of each planar (slots P1-I1, P2-I1, P1-I3, P2-I3, P1-I4, and P2-I4).
13. If three FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) adapters are populated in PHB1, then only FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) may be populated in the remaining slot of that PHB. It must be placed in slots P1-I2 or P2-I2. This allows for six high-performance adapters if the first three in each IO planar are FC 6230.
14. When installing FC 2943: 8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B) in slots P1-I1, P2-I1, P1-I10, and P2-I10, the metal tab on the adapter should be pushed down to assist in plugging.
15. When installing FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) in slot 1, the metal tab on the adapter should be pushed down to assist in plugging.

16. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
17. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
18. Any combination of the following adapters may be installed, but the combined total should not exceed four per planar or LPAR, eight per drawer, or 24 per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
19. The slot usage applies to only FC 6571: I/O Drawer PCI-X Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports.
20. The slot usage applies to only FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports.
21. High performance adapters that are installed in FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports have the following limitations:
 - Two per first PHB on each planar
 - Two per second PHB on each planar
 - One per third PHB on each planar
22. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
23. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7038 pSeries 650

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. Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Figure B-9 show system 7038 6M2 unit rear view with numbered slots.




Figure B-9 System 7038 6M2 unit rear view with numbered slots

Table B-14 on page 847 provides the slot location reference for the 7038 6M2.

Table B-14 Slot location reference 7038 6M2

Slot	PHB	Planar	Loc. Code	Slot Characteristics
1	1	1	Ux.y-P2-I1	64-bit 3.3V, 133 MHz
2	1	1	Ux.y-P2-I2	64-bit 3.3V, 133 MHz
3	1	1	Ux.y-P2-I3	64-bit 3.3V, 133 MHz
4	1	1	Ux.y-P2-I4	64-bit 3.3V, 133 MHz
5	2	1	Ux.y-P2-I5	64-bit 3.3V, 133 MHz
6	2	1	Ux.y-P2-I6	64-bit 3.3V, 133 MHz
7	2	1	Ux.y-P2-I7	32-bit 3.3V, 133 MHz

Note: In Table B-14, Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Table B-15 to identify specific slot location options for the following adapters in your pSeries 650 system.

Table B-15 Placement guidelines 7038-6M2

Feature code	Adapter	Slot usage	Max. per 6M2 Drawer/system	Max. per LPAR	EEH	Hot-plug
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot	3, 5	2/2	1	Y	N
6313	Quad Digital Trunk Telephony PCI Blind Swap, 32/64-bit, 3.3/5V (Type 6313)	1, 2, 3, 4	4/28 (See note 5)	4	Y	Y
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	1, 6, 2, 5, 3, 4, 7	6/32	4	Y	Y (See note 1)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	1, 6, 2, 5, 3, 4, 7	7/15	15	Y	Y (See note 2)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 6, 2, 5, 3, 4	6/30	30	Y	Y

Feature code	Adapter	Slot usage	Max. per 6M2 Drawer/system	Max. per LPAR	EEH	Hot-plug
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	1, 6, 2, 5, 3, 4	6/32 (See note 4)	32	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	6, 2, 5, 3, 4	6/32	32	Y	Y
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	6, 2, 5, 3, 4	4/32	32	Y	Y
6231	SSA (Option for FC 6225/6230)	6, 2, 5, 3, 4	4/32	32	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	6, 2, 5, 3, 4	4/32	32	N/A	N/A
5710	PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter, 32/64-bit, 3.3V (Type 5710)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
5711	PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter, 32/64-bit, 3.3V (Type 5711)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	1, 6, 2, 5, 3, 4	4/30	30	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y

Feature code	Adapter	Slot usage	Max. per 6M2 Drawer/system	Max. per LPAR	EEH	Hot-plug
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 6, 2, 5, 3, 4	6/55	55	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	1, 6, 2, 5, 3, 4	6/30	30	Y	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	1, 6, 2, 5, 3, 4	6/54	14 (See note 3)	Y	Y
4961	10/100 4-Port Ethernet (Type A-E)	1, 6, 2, 5, 3, 4	6/20	20	Y	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	1, 6, 2, 5, 3, 4, 7	7/40	40	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 6, 2, 5, 3, 4, 7	7/8	1	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1, 6, 2, 5, 3, 4, 7	7/8	1	Y	N
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 6, 2, 5, 3, 4, 7	7/30	30	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	1, 6, 2, 5, 3, 4, 7	7/8	1	Y	N
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 6, 2, 5, 3, 4, 7	7/32	32	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 6, 2, 5, 3, 4, 7	7/32	32	Y	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 6, 2, 5, 3, 4, 7	5/53	14 (See note 3)	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. The total of any combination of FC 2947: IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R) and FC 2962: 2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V) must not exceed 14 per LPAR.
4. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
5. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7039 pSeries 655 Model 651

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.

Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed are affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB

from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, all slots on that PHB should be assigned to a single LPAR. For additional information about LPAR (logical partitioning) considerations, refer to the tables in this chapter.

Figure B-10 shows a system 7039 651 unit rear view with numbered slots.




Figure B-10 System 7039 651 unit rear view with numbered slots

Table B-16 provides the slot location reference for the 7039 651.

Table B-16 Slot location reference 7039 651

Slot	PHB	Planar	Loc. Code	Slot Characteristics
1	1	1	Ux.y-P1-I1	64-bit 3.3V, 133 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 133 MHz

Note: In Table B-16, Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Figure B-11 on page 852 shows a system 7039 651 drawer 61D unit rear view with numbered slots.




Figure B-11 System 7039 651 drawer 61D unit rear view with numbered slots

Table B-17 provides the slot location reference for the drawer 61D.

Table B-17 Slot Location Reference drawer 61D

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
1	1	1	Ux.y-P1-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
4	1	1	Ux.y-P1-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
5	2	1	Ux.y-P1-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
6	2	1	Ux.y-P1-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
7	2	1	Ux.y-P1-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	1	Ux.y-P1-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
8	3	1	Ux.y-P1-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
9	3	1	Ux.y-P1-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
10	3	1	Ux.y-P1-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
See note 1	3	1	Ux.y-P1-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
11	1	2	Ux.y-P2-I1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
12	1	2	Ux.y-P2-I2	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz

Slot	PHB	Planar	Loc. code (See note 2)	Slot characteristics (FC 6563)	Slot characteristics (FC 6571)
13	1	2	Ux.y-P2-I3	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
14	1	2	Ux.y-P2-I4	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
15	2	2	Ux.y-P2-I5	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
16	2	2	Ux.y-P2-I6	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
17	2	2	Ux.y-P2-I7	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
See note 1	2	2	Ux.y-P2-Z1	64-bit 3.3V, 33/66 MHz	64-bit 3.3V, 66/133 MHz
18	3	2	Ux.y-P2-I8	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
19	3	2	Ux.y-P2-I9	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
20	3	2	Ux.y-P2-I10	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz
See note 1	3	2	Ux.y-P2-Z2	64-bit 5V, 33 MHz	64-bit 3.3V, 66/133 MHz

Notes:

1. Integrated SCSI adapter.
2. Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position

Use Table B-18 to identify specific slot location options for the adapters in your 7039 pSeries 655 Model 651 system.

Table B-18 Placement guidelines 7039 651

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot	0	0	0	3, 1	0	2	2/2	Y	N

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	20, 10	10, 20	20	0	1	0	1/2	Y	Y
2751	S/390 ESCON Channel, 32-bit, 5V. (Type 5-5)	19, 9, 20, 10	9, 19, 10, 20	19, 20	0	2	0	2/4	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) (See note 7)	11, 1, 13, 3, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17, 19	3, 2, 1	5	3	8/13	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	11, 1, 13, 3, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17, 19	3, 2, 1	5	3	8/13	Y	Y
6230	FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	11, 1, 13, 3, 14, 4, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 4, 14, 5, 15, 7, 17, 9, 19	1, 11, 3, 13, 4, 14, 5, 15, 7, 17, 9, 19	3, 2, 1	6	2	8/14	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	11, 1, 13, 3, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17	2, 1	5	2	7/12	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	11, 1, 13, 3, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17	2, 1	5	2	7/12	Y	Y

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	11, 1, 13, 15, 5, 17, 7, 19, 9	1, 11, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17, 19	2, 1	5	2	7/12	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	11, 1, 13, 15, 5, 17, 7, 19, 9	1, 11, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17, 19	2, 1	5	2	7/12	Y	Y
6203	Dual-Channel Ultra3SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	11, 1, 13, 3, 15, 5, 17, 7, (20 or 19), (10 or 9)	1, 11, 3, 13, 5, 15, 7, 17, (10 or 9), (20 or 19)	11, 13, 15, 17, (20 or 19)	3, 2, 1	5	3	8/13	Y	Y
5710	PCI- X Dual Channel Ultra320 SCSI Blind Swap Adapter, 32/64-bit, 3.3V (Type 5710)	If two FC 6563: 11, 1, 13, 3, 15, 5, 17, 7, 20, 10 If two FC 6571: 11, 1, 15, 5, 18, 8, 12, 2, 16, 6, 19, 9, 13, 3, 17, 7, 20, 10, 14, 4 If mixed planars: 11, 1, 13, 5, 15, 8, 17, 2, 19, 6, 9, 3, 7, 10, 4	If two FC 6563: 1, 11, 3, 13, 5, 15, 7, 17, 10, 20 If two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	If mixed planars or Two FC 6563: 11, 13, 15, 17, 20 If two FC 6571: 11, 15, 18, 12, 16, 19, 13, 17, 20, 14 1, 5, 8, 2, 6, 9, 3, 7, 10, 4	3, 2, 1	5	3	5/10	Y	Y

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
5711	PCI- X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter, 32/64-bit, 3.3V (Type 5711)	If two FC 6563: 11, 1, 13, 3, 15, 5, 17, 7, 20, 10 If two FC 6571: 11, 1, 15, 5, 18, 8, 12, 2, 16, 6, 19, 9, 13, 3, 17, 7, 20, 10, 14, 4 If mixed planars: 11, 1, 13, 5, 15, 8, 17, 2, 19, 6, 9, 3, 7, 10, 4	If two FC 6563: 1, 11, 3, 13, 5, 15, 17, 10, 20 If two FC 6571: 1, 11, 5, 15, 8, 18, 2, 12, 6, 16, 9, 19, 3, 13, 7, 17, 10, 20, 4, 14	If mixed planars or Two FC 6563: 11, 13, 15, 17, 19 If two FC 6571: 11, 15, 18, 12, 6, 16, 19, 13, 17, 20, 14	3, 2, 1	5	3	5/10	Y	Y
2946	PCITURBO WAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	11, 1, 13, 3, 15, 5, 17, 7, 19, 9	1, 11, 3, 13, 5, 15, 7, 17, 9, 19	11, 13, 15, 17, 19	3, 2, 1	5	3	8/13	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	(15, 16, 17); (5, 6, 7); (11, 12, 13, 14); (1, 2, 3, 4); (18, 19, 20); (8, 9, 10) (See note 3)	(5, 6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 3)	(15, 16, 17); (11, 12, 13, 14); (18, 19, 20) (See note 3)	1, 2, 3	2	1	3/5	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) FC 2737: Keyboard/ Mouse Attachment Card - PCI, 32-bit, 3.3/5V	(15, 16, 17); (5, 6, 7); (11, 12, 13, 14); (1, 2, 3, 4); (18, 19, 20); (8, 9, 10) (See note 3)	(5, 6, 7); (15, 16, 17); (1, 2, 3, 4); (11, 12, 13, 14); (8, 9, 10); (18, 19, 20) (See note 3)	(15, 16, 17); (11, 12, 13, 14); (18, 19, 20) (See note 3)	1, 2, 3	2	1	3/5	Y	N
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	11, 1, 12, 2, 13, 3, 14, 4, 15, 5, 16, 6, 17, 7, 18, 8, 19, 9, 20, 10	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	11, 12, 13, 14, 15, 16, 17, 18, 19, 20	0	2	0	2/4	Y	Y

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	11, 1, 12, 2, 13, 3, 14, 4, 15, 5, 16, 6, 17, 7, 18, 8, 19, 9, 20, 10	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	11, 12, 13, 14, 15, 16, 17, 18, 19, 20	0	2	0	2/4	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	11, 1, 12, 2, 13, 3, 14, 4, 15, 5, 16, 6, 17, 7, 18, 8, 19, 9, 20, 10	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	11, 12, 13, 14, 15, 16, 17, 18, 19, 20	1, 2, 3	10	3	13/23	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	11, 1, 12, 2, 13, 3, 14, 4, 15, 5, 16, 6, 17, 7, 18, 8, 19, 9, 20, 10	1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19, 10, 20	11, 12, 13, 14, 15, 16, 17, 18, 19, 20	1, 2, 3	5	3	8/13	Y	Y
2944	128-Port Asynch Controller PCI, 32-bit, 3.3/5V (Type 3-C)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	0	10	0	10/20	Y	Y
2943	8-Port Asynch EIA-232E/R S-422A PCI, 32-bit, 3.3/5V (Type 3-B)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	0	4	0	4/8	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	10	3	13/23	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	10	3	13/23	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	10	3	13/23	Y	Y

Feature code	Adapter	Slot Usage for 61D when there are two 651 and one 61D	Slot Usage for 61D when there is one 651 and one 61D w/FC 3145 (See note 4)	Slot Usage for 61D when there is one 651 and one 61D w/o FC 3145 (See notes 4 and 5)	Slot Usage for one 651	Max. per 61D Planar	Max. per 651	Max. per System with one 651 and one 61D Planar/one 651 and two 61D Planars	E E H	Hot-Plug
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	0	4	0	4/8	N	Y
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	0	4	0	4/8	N	Y
2962	2-Port Multi-protocol PCII, 32-bit, 3.3/5V (Type 9-V)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	0	3	0	3/6	Y	Y
2947	IBM ARTIC 960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	20, 10, 19, 9, 17, 7, 16, 6, 14, 4, 13, 3, 12, 2,	10, 20, 9, 19, 7, 17, 6, 16, 4, 14, 3, 13, 2, 12	20, 19, 17, 16, 14, 13, 12	1, 2, 3	8	3	11/19	Y	Y
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	4	1	5/9	Y	Y (See note 2)
4963	Cryptographic Co-processor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	4	3	7/11	Y	Y (See note 1)
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	20, 10, 19, 9, 18, 8, 17, 7, 16, 6, 15, 5, 14, 4, 13, 3, 12, 2, 11, 1	10, 20, 9, 19, 8, 18, 7, 17, 6, 16, 5, 15, 4, 14, 3, 13, 2, 12, 1, 11	20, 19, 18, 17, 16, 15, 14, 13, 12, 11	1, 2, 3	5	3	8/13	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) and FC 2737: Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V must be placed on the same PHB. All other adapters on that PHB must be on the same LPAR.
4. FC 3145: Remote I/O Cable, 0.5 M is used to connect between the left and right sections of an I/O drawer containing multiple I/O sections. It can also be used to connect two adjacent I/O drawers.
5. This column of the table addresses the case when only one side (slots 11 through 20) of a 61D I/O drawer is connected to a 651 drawer. If more than 10 adapters are present, use both sides of the 61D I/O drawer, and refer to one of the other slot usage columns for the 61D drawer (listed in this table).
6. High performance adapters that are installed in FC 6563: I/O Drawer PCI Planar, 10 Slot, 2 Integrated Ultra3 SCSI Ports have the following limitations:
 - Two per first PHB on each planar
 - Two per second PHB on each planar
 - One per third PHB on each planar
7. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per PCI Host Bridge (PHB) chip is recommended. More than three of these adapters can be installed per PHB, but system performance will not be increased by installing additional adapters to a PHB.
8. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7043 Model 150

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 B-12 on page 860 shows the Model 150 rear view with numbered slots.




Figure B-12 Model 150 system unit rear view with numbered slots

Some adapters must be placed in specific system unit slots to function correctly at optimum performance.

Use Table B-19 to determine where to install an adapter in your system unit. Many of the following notes refer to optimizing system performance.

If you are running AIX, refer to Table B-20 on page 865 to install adapter combinations for your 7043 Model 150. Windows NT (Power PC Edition) does not support more than one graphics adapter in a system unit.

Use the rear-view diagram Figure B-12 to identify slot locations described in Table B-19.

Table B-19 Placement guidelines 7043 150

Feature code	Adapter	Usage	System maximum
4951	10/100 4-Port Ethernet (Type 9-Z)	Slots 2P or 3P	2 (See notes 11 and 13)
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	Slots 1P-5P	4
2920	16 Mbps PCI Token Ring (Type 9-P)	Slots 1P-5P	4
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	Slots 4P, 3P, 2P	3
2948	IBM ARTIC 960 4-Port T1/E1 PCI (Type 9-S)	4P, 3P, 2P	3
2639	Ultimedia Video Capture	Slots 1P-5P	1
2838	POWER GXT120P	Slots 1P-5P	4
2830	POWER GXT130P	Slots 1P-5P	4

Feature code	Adapter	Usage	System maximum
2851	POWER GXT250P	Slots 1P-5P	4
2852	POWER GXT255P	Slots 1P-5P	4
2841	POWER GXT300P	Slots 1P-5P	4
2842	POWER GXT4500P Graphics Adapter (Type 1-Y)	Slots 2P-4P	3
2845	POWER GXT550P	Slots 2P, 3P	1
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	Slots 2P, 3P, 1P, 4P, 5P	4
2849	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	Slots 2P, 3P, 1P, 4P, 5P	4
2823	POWER GXT2000P	Slots 1P-5P	4
2825	POWER GXT3000P	Slot 3P (See note 2)	1
2741	SysKonnnect SK-NET FDDI-LP SAS PCI	Slots 1P-5P (See note 3)	2
2742	SysKonnnect SK-NET FDDI-LP DAS PCI	Slots 1P-5P (See note 3)	2
2743	SysKonnnect SK-NET FDDI-UP SAS PCI	Slots 1P-5P (See note 3)	2
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	2P, 3P	2
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	2P, 3P, 4P, 5P (See notes 7 and 8)	5
6205	PCI Dual Channel Ultra2 SCSI (Type 4-R)	Slots 2P, 3P, 4P, 5P (See notes 7, 8, and 14)	2
6206	PCI Single-Ended Ultra SCSI	Slots 5P, 4P, 3P, 2P, 1P (See note 14)	2
6207	PCI Differential Ultra SCSI	Slots 1P-5P	2
2494	PCI 3-Channel Ultra SCSI RAID (Type 4-T)	Slots 1P-5P (See notes 10 and 15)	2
6215	PCI SSA Multi-Initiator/RAID EL (Type 4-N)	Slots 2P, 3P	1

Feature code	Adapter	Usage	System maximum
6225	Advanced Serial RAID (Type 4-P)	Slots 2P, 3P	1 (See note 6)
6230	Advanced Serial RAID Plus (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	Slots 2P, 3P, 4P	1 (See note 6)
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	Slots 1P-5P	2 (See note 5)
4958	Cryptographic Coprocessor (Type 6-H)	3P, 4P	2 (See note 16)
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	Slots 1P-5P	2 (See note 5)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	Slots 1P-5P	4
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	Slots 1,2	2 (See note 19)
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	Slots 2P, 3P, 1P, 4P, 5P	4 (See note 4)
2968	10/100 Mbps Ethernet PCI	Slots 1P-5P	4 (See note 4)
2943	8-Port Asynchronous EIZ-232E/RS-422A PCI	Slots 1P-5P	2
2944	128-Port Asynchronous 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 (See note 5)
2988	155 Turboways ATM PCI MMF	Slots 1P-5P	2 (See note 5)
2708	ISDN Basic Rate PCI	Slots 1P-5P	1
2998	Turboways 25 ATM PCI	Slots 1P-5P	4

Feature code	Adapter	Usage	System maximum
6310	IBM ARTIC960RxD Quad Digital Trunk PCI (Type 6-E)	Slots 4P, 3P, 2P (See note 9)	3
6311	IBM ARTIC960RxF Quad Digital Trunk PCI (Type 6-G)	Slots 4P, 3P, 2P (See note 9)	2

Notes:

1. The FC 2825: POWER GXT3000P (Type 1-R) occupies the physical space of two adapters in the 7043 Model 150, and, if installed, will fill slots 1 and 2.
2. If both the FC 2838: POWER GXT120P (Type 1-P) or FC 2830: POWER GXT130P (Type 1-T) and the FC 2852: POWER GXT255P (Type 1-N) or FC 2841: POWER GXT300P (Type1-U) are installed, the FC 2852: POWER GXT255P (Type 1-N) or FC 2841: POWER GXT300P (Type1-U) must be installed in slot 2 or 3.
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, place FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) in slots 2 and 3, with no more than two per system. You may have a total of five per system for connectivity, including integrated Ethernet.
5. For optimum system performance, when using a FC 2988: TURBOWAYS 155 PCI MMF ATM (Type9-F), FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), or FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) adapter in LAN Emulation (LANE) a maximum of one adapter per system is recommended.
6. Install only one of the following adapters:
 - FC 6215: PCI SSA Multi-Initiator/RAID EL (Type 4-N)
 - FC 6225: Advanced SSA SerialRAID (Type 4-P)
 - FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)
7. Do not install an FC 6205: PCI Dual Channel Ultra2 SCSI (Type 4-R) or FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) adapter on the same PCI bus as a FC 2838: POWER GXT120P (Type 1-P) or older versions of the FC 2830: POWER GXT130P (Type 1-T). To determine if you have an older version of the FC 2830: POWER GXT130P (Type 1-T), remove the card and examine the bar code label on the back side. If it has a number starting with 11S94H1235, you have an older version of the card.
8. If the FC 6205: PCI Dual Channel Ultra2 SCSI (Type 4-R) or FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y) 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 slots 1, 4 or 5.

9. Do not install more than a combination of three FC 6312: Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312), FC 6311: IBM ARTIC960 RxF PCI, 32-bit, 3.3/5V (Type 6-G), and FC 6310: IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E) adapters in this system.
10. Do not install an FC 2494: PCI 3-Channel Ultra2 SCSI RAID (Type 4-T) and a FC 6311: IBM ARTIC960 RxF PCI, 32-bit, 3.3/5V (Type 6-G) adapter on the same bus.
11. For optimum system performance, use only two 100 Mbps ports per system.
12. Do not install the FC 6205: PCI Dual Channel Ultra2 SCSI (Type 4-R) adapter in slots 2 or 3 if you also have a FC 2823: POWER GXT2000P (Type 1-S) installed on this system.
13. When FC 4951: 10/100 4-Port Ethernet (Type 9-Z) is installed, the following adapters must be installed in slots 1, 4, or 5: FC 2823, FC 2841, FC 2851, and FC 2852.
14. FC 2445 is the internal SCSI cable used with FC 6206: Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K) or with FC 6205: PCI Dual Channel Ultra2 SCSI (Type 4-R). The SCSI adapter card (FC 6206 or FC 6205), which is to be connected to the internal drives using FC 2445, must be installed in Slot 5. If FC 2445 and either FC 6206 or FC 6205 are on the order, slot location placement for FC 6206 and FC 6205 is slot 5 first priority.
15. FC 2494: PCI 3-Channel Ultra2 SCSI RAID (Type 4-T) and FC 2639: Ultimedia Video Capture (Type 7-9) adapters cannot be installed together.
16. Any combination of the following adapters may be installed, but the combined total should not exceed two per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
17. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7043 Model 150 multiple graphics adapter placement guide

Use Table B-20 on page 865 to identify adapter location combination options for these adapters in your 7043 Model 150.

Do not install more than four graphics adapters in your 7043 Model 150.

Table B-20 7043 Model 150 multiple graphics adapter placement table

Adapter	Slot
<p>FC 2838: POWER GXT120P (Type 1-P) or FC 2830: POWER GXT130P (Type 1-T) or FC 2851: POWER GXT250P (Type 1-M) or FC 2852: POWER GXT255P (Type 1-N) or FC 2823: POWER GXT2000P (Type 1-S) or FC 2841: POWER GXT300P (Type1-U) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p> <p>with a</p>	<p>1P -5P</p>
<p>FC 2838: POWER GXT120P (Type 1-P) or FC 2830: POWER GXT130P (Type 1-T) or FC 2851: POWER GXT250P (Type 1-M) or FC 2852: POWER GXT255P (Type 1-N) or FC 2823: POWER GXT2000P (Type 1-S) or FC 2841: POWER GXT300P (Type1-U) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p>	<p>1P-5P</p>

Adapter	Slot
<p>FC 2838: POWER GXT120P (Type 1-P) or FC 2830: POWER GXT130P (Type 1-T) or FC 2851: POWER GXT250P (Type 1-M) or FC 2852: POWER GXT255P (Type 1-N) or FC 2823: POWER GXT2000P (Type 1-S) or FC 2841: POWER GXT300P (Type1-U) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p> <p>with a</p> <p>FC 2855: POWER GXT550P (Type 1-J)</p>	<p>1P</p> <p>2P,3P</p>
<p>FC 2838: POWER GXT120P (Type 1-P) or FC 2830: POWER GXT130P (Type 1-T) or FC 2851: POWER GXT250P (Type 1-M) or FC 2852: POWER GXT255P (Type 1-N) or FC 2823: POWER GXT2000P (Type 1-S) or FC 2841: POWER GXT300P (Type1-U) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p> <p>with a</p> <p>FC 2825: POWER GXT3000P (Type 1-R)</p>	<p>1P or 4P or 5P</p> <p>3P</p>
<p>FC 2639: Ultimedia Video Capture (Type 7-9)</p> <p>with a</p> <p>FC 2825: POWER GXT3000P (Type 1-R)</p>	<p>1P or 4P or 5P</p> <p>3P</p>

Adapter	Slot
FC 2639: Ultimedia Video Capture (Type 7-9) with a FC 2823: POWER GXT2000P (Type 1-S)	1P - 5P 1P - 5P
FC 2639: Ultimedia Video Capture (Type 7-9) with a FC 2841: POWER GXT300P (Type1-U)	1P - 5P 1P - 5P

Notes:

1. The FC 2825: POWER GXT3000P (Type 1-R) occupies the physical space of two adapters in the 7043 Model 150.

7044 Model 170

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 B-13 shows a system 7044 170 unit rear view with numbered slots.




Figure B-13 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 Table B-21 to determine where to install an adapter in your system unit.

Use the rear view diagram in Figure B-13 to identify slot locations described in Table B-21.

Table B-21 Placement guidelines 7044 170

Feature code	Adapter	Slot usage	System maximum
2826	POWER GXT4000P (Type 1-V)	Slot 1, 2, 3, 4, 5, 6	4
2827	POWER GXT6000P (Type 1-W)	Slot 1, 2, 3, 4, 5, 6	2
2843	POWER GXT6500P (Type 1-Z)	Slot 1, 2, 3, 4, 5, 6	4

Feature code	Adapter	Slot usage	System maximum
2842	POWER GXT4500P (Type 1-Y)	Slot 1, 2, 3, 4, 5, 6	4
2825	POWER GXT3000P (Type 1-R)	Slot 2 (See note 11)	1
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	Slot 1, 2, 3, 4, 5, 6	6 (See notes 2 and 6)
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	Slot 1, 2, 3, 4, 5, 6	6 (See notes 2 and 6)
2823	POWER GXT2000P	Slot 1, 2, 3, 4, 5, 6	4
2830	POWER GXT130P	Slot 6, 5, 4, 3	4
2841	POWER GXT300P	Slot 1, 2, 3, 4, 5, 6	4
2849	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	Slot 1, 2, 3, 4, 5, 6	4 (See note 17)
2848	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 1-X)	Slot 1, 2, 3, 4, 5, 6	4 (See note 17)
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Slot 6, 5, 4, 3, 2, 1	3
6311	IBM ARTIC960RxF Quad Digital Trunk PCI (Type 6-G)	Slot 6, 5, 4, 3, 2, 1	2
2494	PCI 3-Channel Ultra SCSI RAID	Slot 6, 5, 4, 3 (See note 1)	2
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	Slot 1, 2, 6, 5, 4, 3	4 (See note 10)
2708	ISDN Basic Rate PCI	Slot 6, 5, 4, 3, 2, 1	1
2639	Ultimedia Video Capture	Slot 6, 5, 4, 3 See note 1	1
6230	SSA Advanced SerialRAID Plus Adapter	Slot 6, 5, 4, 3, 2, 1	2 (See notes 7 and 8)
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	6, 5, 4, 3, 2, 1	4

Feature code	Adapter	Slot usage	System maximum
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	Slot 6, 5, 4, 3, 2, 1	6 (See note 2)
2947	IBM ARTIC960Hx 4-Port Multiprotocol PCI Adapter	Slot 6, 5, 4, 3, 2, 1	2
2943	8-Port Asynchronous EIA-232/RS-422 PCI	Slot 6, 5, 4, 3	2
2944	128-Port Asynchronous Controller PCI	Slot 6, 5, 4, 3	2
2962	2-Port Multiprotocol PCI	Slot 6, 5, 4, 3	2
2963	155 Turboways ATM PCI UTP	Slot 6, 5, 4, 3, 2, 1	6 (See notes 3, 14, and 15)
2985	Ethernet 10base2 PCI (Type 8-Y)	Slot 6, 5, 4,3	4
2987	Ethernet 10base2 PCI (Type 8-Z)	Slot 6, 5, 4,3	4
2988	155 Turboways ATM PCI MMF	Slot 6, 5, 4, 3, 2, 1	6 (See notes 3, 14, and 15)
2742	SysKonnect SK-NET FDDI-LP DAS PCI	Slot 6, 5, 4, 3, 2, 1	6 (See notes 4 and 12)
2741	SysKonnect SK-NET FDDI-LP SAS PCI	Slot 6, 5, 4, 3, 2, 1	6 (See notes 4 and 12)
2743	SysKonnect SK-NET FDDI-UP SAS PCI	Slot 6, 5, 4, 3, 2, 1	6 (See notes 4 and 12)
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	Slot 6, 5, 4, 3, 2, 1	4
6205	PCI Dual Channel Ultra2 SCSI	Slot 6, 5, 4, 3, 2, 1	4
6227	Gigabit Fibre Channel PCI (Type 4-S)	Slot 6, 5, 4, 3, 2, 1	2 (See notes 9 and 20)

Feature code	Adapter	Slot usage	System maximum
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	Slot 6, 5, 4, 3, 1, 2 (See note 18)	2 (See notes 16 and 20)
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	Slot 6, 5, 4, 3, 2, 1	2 (See note 21)
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	Slot 1, 2	2
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	Slot 1, 2	2
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Slot 1, 2, 3, 4, 5, 6	6
5701	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Slot 1, 2, 3, 4, 5, 6	6
4951	10/100 4-Port Ethernet (Type 9-Z)	Slot 6, 5, 4, 3	4 (See note 13)
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	Slot 6, 5, 4, 3, 2, 1	6 (See notes 3, 14, and 15)
4958	Cryptographic Coprocessor (Type 6-H)	Slot 1, 2	2 (See note 19)
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	1, 2	2 (See note 19)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	Slot 6, 5, 4, 3, 2, 1	4
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	Slot 6, 5, 4, 3, 2, 1	2 (See notes 3, 14, and 15)
4959	4/16 Mbps token ring PCI	Slot 6, 5, 4, 3, 2, 1	4

Feature code	Adapter	Slot usage	System maximum
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	Slot 6, 5, 4, 3, 2, 1	5 (See note 13)
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	Slot 6, 5, 4, 3, 2, 1	4 (See note 5)
6206	PCI Single-Ended Ultra SCSI	Slot 6, 5, 4, 3	4
6204	PCI Universal Differential Ultra SCSI	Slot 6, 5, 4, 3, 2, 1	4
2968	10/100 Mbps Ethernet PCI	Slot 6, 5, 4, 3, 2, 1	4 (See note 5)
2998	Turboways 25 ATM PCI	Slot 6, 5, 4, 3, 2, 1	4

Notes:

1. Do not install FC 2494: PCI 3-Channel Ultra2 SCSI RAID (Type 4-T) and FC 2639: Ultimedia Video Capture (Type 7-9) adapters on the same system.
2. For optimum system performance, a combination of the following adapters should not exceed a maximum of one adapter per system:
 - FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)
 - FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
 - FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)

Use slots 1 or 2 if available.
3. For optimum system performance, install up to five FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J), FC 2988: TURBOWAYS 155 PCI MMF ATM (Type 9-F), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), or FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) adapters. Install up to three of these adapters if you are in MTU 1500 mode.
4. For optimum system performance, install up to five SysKconnect adapters in this system.
5. For optimum system performance, install up to three FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters in this system.
6. If your adapter is configured for 100 Mbps mode, install a maximum of two for optimum performance.

7. FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) can be selected with or without the cache upgrade (FC 6235).
8. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230 (1 maximum).
9. A maximum of two FC 6227: Gigabit Fibre Channel PCI (Type 4-S) adapters are supported. One must be in slot 3, 4, 5 or 6; the second in either slot 1 or 2.
10. For FC 2498: PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X) adapter cards shipping in 7044-170, the plastic handle on the end of the card must be replaced by a Metal Handle (PN 09P1971) and two screws (PN 1621170).
11. When placed in slot 2, the FC 2825: POWER GXT3000P (Type 1-R) covers both slots 2 and 3.
12. For optimum system performance, a maximum of five FC 2741: SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *), FC 2742: SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *), and FC 2743: SysKonnnect SK-NET FDDI-UP SAS PCI (Type *) adapters per system is recommended.
13. For optimum system performance, a maximum of one FC 4961: 10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E) or FC 4951: 10/100 4-Port Ethernet (Type 9-Z) adapters per system (with a maximum of three ports) is recommended.
14. For optimum system performance, a maximum of three is recommended for the FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J) and FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C) MTU 1500, and FC 2988: TURBOWAYS 155 PCI MMF ATM (Type 9-F) and FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) MTU 1500 adapters.
15. For optimum system performance, a maximum of five per system is recommended for both the FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J) and FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C) MTU 9180 adapters, and FC 2988: TURBOWAYS 155 PCI MMF ATM (Type 9-F) and FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) MTU 9180 adapters.
16. A maximum of two FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) adapters are supported. One must be in slot 3, 4, 5 or 6; the second in either slot 1 or 2.
17. A maximum of three FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) or FC 2849: POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849) adapters are supported in slots 3 through 6.

18. FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) will operate in 32-bit slots, but at a decrease in performance.
19. Any combination of the following adapters may be installed, but the combined total should not exceed two per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
20. Use of one FC 6227: Gigabit Fibre Channel PCI (Type 4-S) or FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) per bus is recommended. More than one of these adapters can be installed per bus, but system performance will not be increased by installing additional adapters to a bus.
21. A maximum of two FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) adapters are supported. One must be in slot 3, 4, 5 or 6; the second in either slot 1 or 2.
22. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7044 Model 170 multiple graphics adapter placement guide

Use Table B-22 to identify adapter location combination options for these adapters in your 7044 Model 170

Do not install more than four graphics adapters in your 7044 Model 170.

Table B-22 7044 Model 170 multiple graphics adapter placement table

Adapter	Slots
FC 2823: POWER GXT2000P (Type 1-S) or FC 2841: POWER GXT300P (Type1-U) with a	1-6
FC 2825: POWER GXT3000P (Type 1-R) or FC 2827: POWER GXT6000P (Type 1-W) or FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z)	2 - 4 (See note)

Adapter	Slots
<p>FC 2827: POWER GXT6000P (Type 1-W) or FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z)</p> <p>with a</p> <p>FC 2825: POWER GXT3000P (Type 1-R) or FC 2827: POWER GXT6000P (Type 1-W) or FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z) or FC 2826: POWER GXT4000P (Type 1-V) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p>	<p>2 (See note)</p> <p>4</p>
<p>FC 2639: Ultimedia Video Capture (Type 7-9)</p> <p>with a</p> <p>FC 2825: POWER GXT3000P (Type 1-R) or FC 2827: POWER GXT6000P (Type 1-W) or FC 2843: POWER GXT6500P Graphics Adapter (Type 1-Z)</p>	<p>3, 4, 5, 6</p> <p>2 (See note)</p>
<p>FC 2639: Ultimedia Video Capture (Type 7-9)</p> <p>with a</p> <p>FC 2823: POWER GXT2000P (Type 1-S) or FC 2826: POWER GXT4000P (Type 1-V) or FC 2842: POWER GXT4500P Graphics Adapter (Type 1-Y)</p>	<p>3, 4, 5, 6</p> <p>1 or 2 (See note)</p>
<p>FC 2639: Ultimedia Video Capture (Type 7-9)</p> <p>with a</p> <p>FC 2841: POWER GXT300P (Type1-U)</p>	<p>3 - 6</p> <p>1 - 6</p>

Adapter	Slots
FC 2830: POWER GXT130P (Type 1-T) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) with a FC 2823: POWER GXT2000P (Type 1-S)	3 -6 1 or 2
FC 2830: POWER GXT130P (Type 1-T) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) with a FC 2825: POWER GXT3000P (Type 1-R)	3, 4, 5, 6 2 (See note)
FC 2830: POWER GXT130P (Type 1-T) or FC 2848: GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X) with a FC 2841: POWER GXT300P (Type1-U) or FC 2639: Ultimedia Video Capture (Type 7-9)	3 -6 3 -6

Note:

Since slots 1 and 2 are capable of running at 50 MHz, the FC 2823: POWER GXT2000P (Type 1-S) and FC 2825: POWER GXT3000P (Type 1-R) should be placed in these slots for optimum performance. Take care when placing any 33 Mhz PCI adapter in slots 1 and 2, because this forces the slots to run at 33 Mhz.

7044 Model 270

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 B-14 on page 877 shows a system 7044 270 unit rear view with numbered slots.




Figure B-14 Model 270 system unit rear view with numbered slots

Use Table B-23 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 B-14 to identify slot locations described in Table B-23.

Table B-23 Placement guidelines 7044 270

Feature Code	Adapter	Slot Usage	System Maximum
2826	POWER GXT4000P (Type 1-V)	Slot 1, 2, 3, 4, 5	4
2843	POWER GXT6500P (Type 1-Z)	Slot 1, 2, 3, 4, 5	2
2842	POWER GXT4500P (Type 1-Y)	Slot 1, 2, 3, 4, 5	4
2827	POWER GXT6000P (Type 1-W)	Slot 2, 4	2

Feature Code	Adapter	Slot Usage	System Maximum
2825	POWER GXT3000P (See note 8)	Slot 2, 4	2
2848	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 1-X)	Slot 1, 2, 3, 4, 5	4
2849	GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	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
2823	POWER GXT2000P	Slot 1, 2, 3, 4, 5	4
2838	POWER GXT120P	Slot 1, 2, 3, 4, 5	4
2830	POWER GXT130P	Slot 5, 4, 3	3
2841	POWER GXT300P	Slot 1, 2, 3, 4, 5	4
2969	Gigabit Ethernet-SX PCI	Slot 1, 2, 3, 4, 5 (See notes 2 and 4)	5
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	Slot 1, 2, 3, 4, 5 (See notes 2 and 4)	5
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Slot 1, 2, 3, 4, 5	5
5701	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	Slot 1, 2, 3, 4, 5	5
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	Slot 1, 2	2
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	1, 2, 3, 4, 5	2
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	Slot 1, 2	2
6310	IBM ARTIC960RxD Quad Digital Trunk PCI	Slot 5, 4, 3, 2, 1	3
6311	IBM ARTIC960RxF Quad Digital Trunk PCI (Type 6-G)	Slot 5, 4, 3, 2, 1	3
2494	PCI 3-Channel Ultra SCSI RAID	Slot 5, 4, 3	2 (See note 1)

Feature Code	Adapter	Slot Usage	System Maximum
2708	ISDN Basic Rate PCI	Slot 5, 4, 3	1
2639	Ultimedia Video Capture	Slot 5, 4, 3	1 (See note 1)
6215	PCI SSA Multi-Initiator/RAID EL	Slot 3, 4, 5	1
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	Slot 1, 2, 3, 4, 5	5 (See note 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 Asynchronous Controller PCI	Slot 1, 2, 3, 4, 5	2
2962	2-Port Multiprotocol PCI	Slot 5, 4, 3	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
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	Slot 1, 2, 3, 4, 5	5
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	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	5
4951	10/100 4-Port Ethernet (Type 9-Z)	Slot 5, 4, 3	3 (See note 5)
4959	4/16 Mbps token ring PCI	Slot 1, 2, 3, 4, 5	4
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	Slot 1, 2, 3, 4, 5	5 (See note 5)
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	Slot 1, 2, 3, 4, 5	5 (See note 3)

Feature Code	Adapter	Slot Usage	System Maximum
4958	Cryptographic Coprocessor (Type 6-H)	Slot 2, 1	2 (See note11)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	Slot 5, 4, 3, 2, 1	4
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-l)	Slot 2, 1	2 (See note 11)
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	Slot 1, 2, 3, 4, 5	2
6205	PCI Dual Channel Ultra2 SCSI	Slot 1, 2, 3, 4, 5	2
6206	PCI Single-Ended Ultra SCSI	Slot 5, 4, 3	2
6207	PCI Differential Ultra SCSI	Slot 1, 2, 3, 4, 5	4
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	Slot 1, 2, 3, 4, 5	4
6227	Gigabit Fibre Channel PCI (Type 4-S)	Any slot (See notes 9 and 12)	2
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	Slot 2, 1 (See notes 10 and 12)	2
6230	SSA Advanced SerialRAID Plus Adapter FC 6230+6231: 4-Port SSA 40 with 128MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	Slot 1, 2, 3, 4, 5	2 (See notes 6 and 7)
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	Slot 1, 2, 3, 4, 5	2
2968	10/100 Mbps Ethernet PCI	Slot 1, 2, 3, 4, 5	5 (See note 3)
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 (See note 2)

Notes:

1. Do not install FC 2639: Ultimedia Video Capture (Type 7-9) and FC 2494: PCI 3-Channel Ultra2 SCSI RAID (Type 4-T) adapters on the same system.
2. For optimum system performance, a combination of the following adapters should not exceed a maximum of two adapters per system:
 - FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)
 - FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
 - FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)

Install them in slots 1 and 2 (if available).

3. 7044 Model 270 supports up to five FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters per system, making the total six if you use the Integrated Ethernet port.
4. For optimum system performance, install a maximum combination of two FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A) adapters per system, or two FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U), or one of each, when configured to run at 1000 Mbps. If the adapter is configured for 100 Mbps mode, use up to five of these adapters (plus the integrated 100 Mbps port) for optimum performance.
5. For optimum system performance, install a maximum of two FC 4951: 10/100 4-Port Ethernet (Type 9-Z) or FC 4961: 10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E) adapters in 100 Mbps mode per system.
6. FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) can be selected with or without the cache upgrade (FC 6235: Fast Write Cache Option for FC 6225 and FC 6230).
7. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a cache upgrade for the FC 6225: Advanced SSA SerialRAID (Type 4-P) adapter and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) (maximum of one per 6230).
8. If two FC 2825: POWER GXT3000P (Type 1-R) adapters are placed in slots 2 and 4, then slot 3 is covered and cannot be used.
9. FC 6227: Gigabit Fibre Channel PCI (Type 4-S) and FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) adapters are restricted to only one adapter in slot 1 or 2 and/or one adapter in slot 3 or 4 or 5.
10. FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) will operate in 32-bit slots, but at a decrease in performance.

- FC 2639: Ultimedia Video Capture (Type 7-9) cannot be placed in slot 1.

7046 Model B50

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 B-15 shows a system 7046 B50 unit rear view with numbered slots.




Figure B-15 Model B50 system unit rear view with numbered slots

Table B-25 summarizes the adapters on the RS/6000 7046 Model B50 system and describes the adapter placement guidelines. Please note that PCI slot1 can hold short cards only.

Table B-25 Placement guidelines 7046 B50

Feature code	Adapter	Slot usage	System maximum
2830	POWER GXT130P (Type 1-T)	1, 2	1
2841	Mirage 2D(GTX300P)	1, 2	1
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1	1
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1	1
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	1, 2	2
6207	PCI Differential Ultra SCSI (Type 4-L)	1, 2	2
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	1, 2	1
6205	PCI Dual Channel Ultra2 SCSI (Type 4-R)	2	1

Feature code	Adapter	Slot usage	System maximum
2494	PCI 3-Channel Ultra2 SCSI RAID (Type 4-T)	2	1
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	2	1
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	2	1
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	2	1
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	N/A	1
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	2	1
2948	IBM ARTIC960 4-Port T1/E1 PCI (Type 9-S)	2	1
2968	10/100 Mbps Ethernet Tx PCI (Type 9-P)	1, 2	2 (See note 1)
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 2	2 (See note 1)
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 2	2
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 2	2
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 2	2
2963	TURBOWAYS 155 PCI UTP ATM (Type 9-J)	1, 2	2 (See note 2)
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	1, 2	2 (See note 2)

Feature code	Adapter	Slot usage	System maximum
2988	TURBOWAYS 155 PCI MMF ATM (Type 9-F)	1, 2	2 (See note 2)
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	1, 2	2 (See note 2)
2708	ISDN Basic Rate PCI (Type 9-N)	1, 2	2
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	1, 2	1
2741	SysKonnnect SK-NET FDDI-LP SAS PCI, 32-bit, 3.3/5V (Type *)	1, 2	1
2743	SysKonnnect SK-NET FDDI-UP SAS PCI (Type *)	1, 2	1
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 2	2
4951	10/100 4-Port Ethernet (Type 9-Z)	2	2 (See note 3)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	2	1
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	2	1
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	2	1 (See note 4)
4958	Cryptographic Coprocessor (Type 6-H)	2	1 (See note 4)
6225	Advanced SSA SerialRAID (Type 4-P)	2	1

Notes:

1. For optimum system performance, if the FC 2968: 10/100 Mbps Ethernet Tx PCI (Type 9-P) or FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) is used in 100 Mbps mode, then a maximum of one adapter is recommended per system unit.

2. For optimum system performance, when using a FC 2988: TURBOWAYS 155 PCI MMF ATM (Type 9-F), FC 2963: TURBOWAYS 155 PCI UTP ATM (Type 9-J), FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), or FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) adapter in LAN Emulation (LANE), a maximum of one adapter per system is recommended.
3. For optimum system performance, use only two 100 Mbps ports per system.
4. Only one of the following adapters may be installed per system: FC 4958: Cryptographic Coprocessor (Type 6-H) and FC 4963: Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I), also referred to as RPQ 8A1162: Cryptographic Coprocessor.
5. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7311 Model D10

Adapter cards for the 7311 Model D10 plug into PCI adapter slots in the I/O subsystem. Each drawer is capable of handling up to six PCI adapters. All slots in the drawer are capable of using PCI-X adapters. Some adapters must be placed in specific slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed will be affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, then all slots on that PHB should be assigned

to a single LPAR. Refer to the tables in this chapter for additional information about LPAR (logical partitioning) considerations.

Note: For information about the maximum number of adapters per LPAR, refer to the PCI adapter placement information for your system.

Figure B-16 shows a system 7311 D10 unit rear view with numbered slots.




Figure B-16 System 7311 D10 unit rear view with numbered slots

Table B-26 shows the slot location reference for a 7311 D10

Table B-26 Slot location reference 7311 D10

Slot	PHB	Planar	Loc. Code	Slot Characteristics
1	1	1	Ux.y-P1-I1	64-bit 5 V, 33 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 133 MHz
4	2	1	Ux.y-P1-I4	64-bit 3.3V, 133 MHz
5	2	1	Ux.y-P1-I5	64-bit 3.3V, 133 MHz
6	2	1	Ux.y-P1-I6	64-bit 3.3V, 133 MHz

Notes:

1. In Table 54, Ux.y represents the Hardware Management Console (HMC) location code, where x is the rack location and y is the drawer position.
2. Slot 1 is a PCI slot. Slots 2 through 6 are PCI-X slots.

Use Table B-27 to identify specific slot location options for the adapters in your 7311 Model D10 drawer.

Table B-27 Placement guidelines 7311 D10

Feature code	Adapter	Slot Usage	Max. per Drawer	EEH	Hot-Plug
2751	S/390 ESCON Channel, 32-bit, 5V, (Type 5-5)	1	1	Y	Y
6313	Quad Digital Trunk Telephony PCI Blind Swap, 32/64-bit, 3.3/5V (Type 6313)	(1, 2, 3), (4, 5, 6) (See note 3)	3	Y	N
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	1	1	Y	Y
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	1, 6, 2, 5, 3, 4	6	Y	Y (See note 1)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	1, 6, 2, 5, 3, 4	6	Y	Y (See note 2)
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	2, 4, 3, 5, 6, 1	6	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	2, 4, 3, 5, 6, 1	6	Y	Y
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	2, 4, 3, 5, 6, 1	6	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	2, 4, 3, 5, 6, 1	6	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	2, 4, 3, 5, 6, 1	6 (See note 5)	Y	Y
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	2, 4, 3, 5, 6, 1	6 (See note 5)	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	2, 4, 3, 5, 6, 1	6	Y	Y
5710	PCI-X Dual Channel Ultra320 SCSI Blind Swap Adapter, 32/64-bit, 3.3V (Type 5710)	2, 4, 3, 5, 6, 1	6	Y	Y
5711	PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap Adapter, 32/64-bit, 3.3V (Type 5711)	2, 4, 3, 5, 6, 1	6	Y	Y

Feature code	Adapter	Slot Usage	Max. per Drawer	EEH	Hot-Plug
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	2, 4, 3, 5, 6, 1	4	Y	Y
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	N/A	4	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	N/A	4	N/A	N/A
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	2, 4, 3, 5, 6, 1	4	Y	Y
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	1, 6, 2, 5, 3, 4	6	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 6, 2, 5, 3, 4	6	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1, 6, 2, 5, 3, 4	6	Y	N
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	2, 4, 3, 5, 6, 1	6	Y	Y
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	2, 4, 3, 5, 6, 1	2	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	2, 4, 3, 5, 6, 1	6	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	2, 4, 3, 5, 6, 1	6	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	2, 4, 3, 5, 6, 1	6	Y	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	1, 2, 5, 3, 4, 6	6	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 6, 2, 5, 3, 4	6	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	2, 4, 3, 5, 6, 1	6	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	2, 4, 3, 5, 6, 1	6	Y	N

Feature code	Adapter	Slot Usage	Max. per Drawer	EEH	Hot-Plug
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 6, 2, 5, 3, 4	6	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 6, 2, 5, 3, 4	6	Y	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 6, 2, 5, 3, 4	6	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. The slot usage for FC 6313: Quad Digital Trunk Telephony PCI Blind Swap, 32/64-bit, 3.3/5V (Type 6313) is valid when FC 2877: Quad DTA, H.100, 4-Drop Cable is used.
4. The slot usage for FC 6313: Quad Digital Trunk Telephony PCI Blind Swap, 32/64-bit, 3.3/5V (Type 6313) is valid when FC 2877: Quad DTA, H.100, 4-Drop Cable is not used.
5. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) or FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) per bus is recommended. More than one of these adapters can be installed per bus, but system performance may not be increased by installing additional adapters to a bus.
6. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

7311 Model D20

Adapter cards for the 7311 Model D20 plug into PCI adapter slots in the in the I/O subsystem. Each drawer is capable of handling up to seven PCI adapters. All slots in the drawer are capable of using PCI-X adapters.

Some adapters must be placed in specific slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Logical partition (LPAR) considerations

Place redundant devices in separate I/O drawers for best availability performance. Place non-redundant devices in the same I/O drawer. If you place non-redundant devices in one drawer, the system is less exposed to other-drawer failures.

Some devices do not have enhanced error handling (EEH) capabilities built in to their device drivers. If these devices fail, the PCI Host Bridge (PHB) in which they are placed will be affected. If the I/O subsystem encounters a severe error, all slots in the PHB are also affected. To clear this condition, you may reboot the system. In addition, it is also possible to remove the failed PCI slots on an affected PHB from the partition profile or profiles that include these PCI slots, and reboot the partition or partitions that terminated at the time of the error.

To avoid PHB errors related to non-EEH adapters, it is strongly recommended that if a non-EEH adapter is used, then all slots on that PHB should be assigned to a single LPAR. Refer to the tables in this chapter for additional information about LPAR (logical partitioning) considerations.

Note: For information about the maximum number of adapters per LPAR, refer to the PCI adapter placement information for your system.

Figure B-17 on page 893 shows a system 7311 D20 unit rear view with numbered slots.




Figure B-17 System 7311 D20 unit rear view with numbered slots

Table B-28 shows the slot location reference for a 7311 D20.

Table B-28 Slot location reference 7311 D20

Slot	PHB	Planar	Loc. Code	Slot Characteristics
1	1	1	Ux.y-P1-I1	64-bit 5 V, 33 MHz
2	1	1	Ux.y-P1-I2	64-bit 3.3V, 133 MHz
3	1	1	Ux.y-P1-I3	64-bit 3.3V, 133 MHz
4	1	1	Ux.y-P1-I4	64-bit 3.3V, 133 MHz
5	2	1	Ux.y-P1-I5	64-bit 3.3V, 133 MHz
6	2	1	Ux.y-P1-I6	64-bit 3.3V, 133 MHz
7	2	1	Ux.y-P1-I7	64-bit 3.3V, 133 MHz

Note: In Table 56, Ux.y represents the Hardware Management Console (HMC) location code where x is the rack location and y is the drawer position.

Use Table B-29 on page 894 to identify specific slot location options for the adapters in your 7311 Model D20 drawer.

Table B-29 Placement guidelines 7311 D20

Feature code	Adapter	Slot usage	Max. per drawer	EEH	Hot-plug
6312	Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312)	(1, 2, 3, 4); (5, 6, 7) (See note 3)	4	Y	N
8398	SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1 slot	3, 5	2	Y	N
2732	Serial HIPPI, Short-Wave Optics (Type 9-W)	3, 6, 1, 5, 2	2	Y	Y
2733	Serial HIPPI, Long-Wave Optics (Type 9-W)	3, 6, 1, 5, 2	2	Y	Y
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	7, 4, 1, 5, 2, 6, 3	7	Y	Y
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	7, 4, 1, 5, 2, 6, 3	7	Y	Y
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	7, 4, 5, 2, 6, 3, 1	4	Y	Y
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	7, 4, 5, 2, 6, 3, 1	7	Y	Y
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	1, 5, 2, 6, 3, 7, 4	7	N	N
4963	Cryptographic Coprocessor, FIPS-4, 32-bit, 3.3/5V (Type 6-I)	1, 5, 2, 6, 3, 7, 4	7	Y	Y (See note 1)
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	1, 5, 2, 6, 3, 7, 4	7	Y	Y (See note 2)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 5, 2, 6, 3, 7, 4	7	Y	Y
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 5, 2, 6, 3, 7, 4	7	Y	Y
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 5, 2, 6, 3, 7, 4	7	Y	Y

Feature code	Adapter	Slot usage	Max. per drawer	EEH	Hot-plug
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
6228	2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)	1, 5, 2, 6, 3, 7, 4	4	Y	Y
6231	128 MB DRAM Option Card for FC 6225 and FC 6230	N/A	4	N/A	N/A
6235	Fast Write Cache Option for FC 6225 and FC 6230	N/A	4	N/A	N/A
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 5, 2, 6, 3, 7, 4	7	Y	N
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1, 5, 2, 6, 3, 7, 4	7	Y	N
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2737	Keyboard/Mouse Attachment Card - PCI, 32-bit, 3.3/5V	1, 5, 2, 6, 3, 7, 4	7	Y	N

Feature code	Adapter	Slot usage	Max. per drawer	EEH	Hot-plug
2742	SysKonnnect SK-NET FDDI-LP DAS PCI, 32-bit, 3.3/5V (Type *)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 5, 2, 6, 3, 7, 4	7	Y	Y
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 5, 2, 6, 3, 7, 4	7	Y	Y

Notes:

1. Before hot-plugging this adapter, see the *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235 for the required procedures.
2. Before hot-plugging this adapter, see the *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254 for the required procedures.
3. The slot usage for FC 6312: Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312) is valid when FC 4353: Internal H100 BUS, 8-Position Cable or FC 2877: Quad DTA, H.100, 4-Drop Cable is used. FC 2877 and FC 4353 can daisy chain up to four FC 6312 that are adjacent to each other. All FC 6312 will be placed adjacent to each other and if either one FC 2877 or one FC 4353 is on the order, all FC 6312 will be connected together with a single cable
4. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) or FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) per bus is recommended. More than one of these adapters can be installed per bus, but system performance may not be increased by installing additional adapters to a bus.
5. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

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, customers and service representatives can refer to *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280.

The Models 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.

A minimum of one Ethernet adapter is required for Models S70, S7A, and S80 systems to function as an attached servers within the IBM RS/6000 SP environment.

9112-265 IntelliStation POWER 265

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 B-18 shows a system 9112 Model 265 unit rear view with numbered slots.




Figure B-18 System 9112 Model 265 unit rear view with numbered slots

Some adapters must be placed in specific system unit slots to function correctly at optimum performance.

Use the Table B-30 to determine where to install an adapter in your system unit.

Many of the following notes refer to optimizing system performance. Read “System performance” on page 787 for more performance-related information. Use the rear-view diagram in Figure B-18 to identify slot locations described in Table B-30.

Table B-30 Placement guidelines 9112 Model 265

Feature code	Adapter	Slot usage	System maximum
2842	POWER GXT4500P Graphics Adapter (Type 1-Y)	1, 3, 2, 4, 5	2
2843	POWER GXT6500P Graphics Adapter (Type 1-Z)	1, 3, 2, 4, 5	2

Feature code	Adapter	Slot usage	System maximum
6228	Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)	1, 3, 2, 4, 5	2 (See notes 8 and 9)
6239	Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	1, 3, 2, 4, 5	2 (See notes 8 and 9)
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 2	2
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 2	2
2969	Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)	1, 2, 3, 4, 5	5 (See note 4)
2975	Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)	1, 2, 3, 4, 5	5 (See note 4)
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 2, 3, 4, 5	5
8244	Audio PCI Adapter for Workstations (Type 8244)	1, 2, 3, 4, 5	1
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 2, 3, 4, 5	4
2848	GXT135P Graphics Adapter, 32-bit, 3.3/5V (Type 1-X)	1, 2, 3, 4, 5	4
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	5, 4, 3, 2, 1	2
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	1, 2, 3, 4, 5	5
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	1, 2, 3, 4, 5	5
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	1, 2, 3, 4, 5	5
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	1, 2, 3, 4, 5	5

Feature code	Adapter	Slot usage	System maximum
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	1, 2, 3, 4, 5	2 (See notes 1 and 2)
6206	Single-Ended Ultra SCSI Adapter, 32-bit, 5V (Type 4-K)	5, 4, 3	2
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	5, 4, 3, 2, 1	4
4961	10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E)	5, 4, 3, 2, 1	5 (See notes 5 and 7)
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	5, 4, 3, 2, 1	5
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	5, 4, 3, 2, 1	5
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	5, 4, 3, 2, 1	4
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	5, 4, 3, 2, 1	5 (See note 6)
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	5, 4, 3, 2, 1	5
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	5, 4, 3, 2, 1	5 (See note 5)
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	5, 4, 3, 2, 1	5 (See note 5)
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 2, 3, 4, 5	5 (See note 4)

Notes:

1. FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) or FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM can be selected with or without the cache upgrade (FC 6235: Fast Write Cache Option for FC 6225 and FC 6230).
2. FC 6235: Fast Write Cache Option for FC 6225 and FC 6230 is a cache upgrade for the FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM adapter

and can only be selected along with FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) or FC 6230+6231: 4-Port SSA 40 with 128 MB DIMM (maximum of one per FC 6230/FC 6230+6231).

3. FC 6227: Gigabit Fibre Channel PCI (Type 4-S) is restricted to only one adapter in slot 1 or 2 and/or one adapter in slot 3, 4, or 5.
4. For optimum system performance, a combination of the following adapters should not exceed a maximum of one adapter per system:
 - FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)
 - FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
 - FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)Use 64-bit slots (1, 2, or 3) if available.
5. For optimum system performance, install up to 5 FC 4953: 64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C), or FC 4957: 64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D) adapters. Install up to three of these adapters if you are in MTU 1500 mode. Use 64-bit slots (1, 2, or 3) if available.
6. For optimum system performance, install up to three FC 4962: Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F) adapters in this system.
7. For optimum system performance, a maximum of one FC 4961: 10/100 4-Port Ethernet, 32/64-bit, 3.3/5V (Type A-E) adapter per system (with a maximum of three ports) is recommended.
8. FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) and FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) will operate in 32-bit slots, but at a decrease in performance.
9. Use of one FC 6228: 2 Gigabit Fibre Channel Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W) or FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) per bus is recommended. More than one of these adapters can be installed per bus, but system performance will not be increased by installing additional adapters to a bus.
10. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system.

9114-275 IntelliStation POWER 275

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.

Some adapters must be placed in specific system unit slots to function correctly at optimum performance. Use the information in the following sections of this chapter to determine where to install adapters in your system unit.

Figure B-19 shows a system 9114 Model 275 unit rear view with numbered slots.




Figure B-19 System 9114 Model 275 unit rear view with numbered slots

All slots in the system are capable of using PCI-X adapters. Slots 1, 4, 5, and 6 support full-length PCI-X adapters and are 64-bit capable at 133 MHz, 3.3 volts. Slots 2 and 3 support half-length PCI-X adapters and are 32-bit capable at 66 MHz, 3.3 volts.

Use Table B-31 on page 903 to identify specific slot location options for the adapters in your 9114 Model 275 system.

Table B-31 Placement guidelines 9114 Model 275

Feature code	Adapter	Slot usage	System maximum
5709	Dual Channel SCSI RAID Enablement Card (Type 5709)	N/A	1
2843	POWER GXT6500P Graphics Adapter (Type 1-Z)	4	1
2842	POWER GXT4500P Graphics Adapter (Type 1-Y)	4, 5, 6	2
6312	Quad Digital Trunk Telephony PCI, 32/64-bit, 3.3/5V (Type 6312)	5, 6	2
6310	IBM ARTIC960 RxD Quad Digital Trunk PCI, 32-bit, 3.3/5V (Type 6-E)	5, 6	2
6230	Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P) FC 6231: 128 MB DRAM Option Card for FC 6225 and FC 6230 FC 6235: Fast Write Cache Option for FC 6225 and FC 6230	5, 6	2
5712	Dual Channel Ultra320 SCSI 32/64-bit, 3.3V (Type 5712)	5, 6	2
2498	PCI 4-Channel Ultra3 SCSI RAID, 32/64-bit, 3.3/5V (Type 4-X)	5, 6	2
6203	Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)	5, 6	2
5703	PCI-X Dual Channel Ultra320 SCSI RAID Adapter, 32/64-bit, 3.3V (Type 5703)	5, 6	2
5701	10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V	1, 5, 6	3
5700	Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V	1, 5, 6	3
2946	PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)	1, 5, 6	3
5707	2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)	1, 5, 6	3
5706	2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)	1, 5, 6	3
6239	2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)	1, 5, 6	3 (See note 1)
8244	Audio PCI Adapter for Workstations (Type 8244)	2, 3	1
4960	IBM Cryptographic Accelerator, 32-bit, 3.3/5V (Type 6-J)	1, 2, 3, 5, 6	4
6204	PCI Universal Differential Ultra SCSI, 32-bit, 3.3/5V (Type 4-U)	1, 2, 3, 5, 6	4
2849	POWER GXT135P Graphics Accelerator, 32-bit, 3.3/5V (Type 2849)	1, 2, 3, 5, 6	2
2947	IBM ARTIC960Hx 4-Port Selectable PCI, 32-bit, 3.3/5V (Type 9-R)	5, 6	2

Feature code	Adapter	Slot usage	System maximum
4953	64-bit/66 MHz PCI ATM 155 UTP, 3.3/5V (Type A-C)	1, 5, 6	3
4957	64-bit/66 MHz PCI ATM 155 MMF, 3.3/5V (Type A-D)	1, 5, 6	3
4962	Ethernet/LAN Encryption 10/100BaseT, 32-bit, 3.3/5V (Type A-F)	1, 2, 3, 5, 6	4
4959	High-Speed Token Ring PCI, 32-bit, 3.3/5V (Type 9-Y)	1, 2, 3, 5, 6	4
2943	8-Port Asynchronous EIA-232E/RS-422A PCI, 32-bit, 3.3/5V (Type 3-B)	1, 2, 3, 5, 6	4
2944	128-Port Asynchronous Controller PCI, 32-bit, 3.3/5V (Type 3-C)	1, 2, 3, 5, 6	4
2962	2-Port Multiprotocol PCII, 32-bit, 3.3/5V (Type 9-V)	1, 2, 3, 5, 6	4

Notes:


1. Use of one FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704) per bus is recommended. More than one of these adapters can be installed per bus, but system performance may not be increased by installing additional adapters to a bus.
2. For optimum system performance, the combined maximum of high-performance adapters should not exceed the maximums listed for FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706) or FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707) if either of these adapters is installed in the system. For a list of high-performance adapters, see "High-performance adapters" on page 904.

High-performance adapters

The following is a list of high-performance adapters. To improve system performance, combinations of these adapters may be limited to a lower maximum number than the maximum number permitted for individual adapters used in a system or drawer. For additional information about these limitations, see the Notes section of the placement rules for each system.

- ▶ FC 2946: PCI TURBOWAYS 622 PCI MMF ATM, 64-bit, 3.3/5V (Type A-B)
- ▶ FC 2969: Gigabit Ethernet Fibre, 1000BaseT, 64-bit, 3.3/5V (Type 9-U)
- ▶ FC 2975: Gigabit Ethernet (UTP) 1000BaseT, 64-bit, 3.3/5V (Type A-A)
- ▶ FC 5700: Gigabit Ethernet, 1000 Base-SX, 32/64-bit, 3.3/5V
- ▶ FC 5701: 10/100/1000 Base-TX Ethernet, 32/64-bit, 3.3/5V

- ▶ FC 5706: 2-Port 10/100/1000 Base-TX Ethernet PCI-X, 32/64-bit, 3.3/5V (Type 5706)
- ▶ FC 5707: 2-Port Gigabit Ethernet-SX PCI-X, 32/64-bit, 3.3/5V (Type 5707)
- ▶ FC 6203: Dual-Channel Ultra3 SCSI, 32/64-bit, 3.3/5V (Type 4-Y)
- ▶ FC 6228: 2 Gigabit FC Adapter for 32/64-bit, 3.3/5V PCI Bus (Type 4-W)
- ▶ FC 6230: Advanced SerialRAID Plus, 32-bit, 3.3/5V (Type 4-P)
- ▶ FC 6239: 2 Gigabit Fibre Channel Adapter, 32/64-bit, 3.3V (Type 5704)
- ▶ FC 8396: SP System Attachment (Type 6-F)
- ▶ FC 8397: SP System Attachment, 64-bit, 3.3/5V, 2-slot
- ▶ FC 8398: SP Switch2 PCI-X Attachment, 64-bit, 3.3V, 1-slot



Power cord features

This appendix provides the information related to the external power cords available for the following RS/6000, pSeries, and IntelliStation products:

- ▶ 7026 Model B80 (p640)
- ▶ 7028 Model 6C1, and 6E1 (p610)
- ▶ 7028 Model 6C4, and 6E4 (p630)
- ▶ 7029 Model 6C3, and 6E3 (p615)
- ▶ 7038 Model 6M2 (p650)
- ▶ 7039 Model 651 (p655)
- ▶ 7040 Model 671 (p670)
- ▶ 7040 Model 681 (p690)
- ▶ 7043 Model 150 (43p)
- ▶ 7044 Model 170 (44p)
- ▶ 9112 Model 265 (IntelliStation)
- ▶ 9114 Model 275 (IntelliStation)

Table C-1 on page 908 lists the features codes of the common power cords for Model B80, 6C1, 6E1, 6C4, 6E4, 6C3, and 6E3.

Table C-1 Power cords list by feature code

Feature code	Description	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
9800	Power cord - United States/Canada, L6-30P, 250V, single phase, 14 ft., 30A	X	X	X	X	X	X	X
9802	Power cord - Brazil (125V, 15A)	X	X	X	X	X	X	X
9820	Power cord - Belgium, Finland, France (250V, 16A)	X	X	X	X	X	X	X
9821	Power cord - Denmark (250V, 10A)	X	X	X	X	X	X	X
9825	Power cord - United Kingdom (250V, 13A)	X	X	X	X	X	X	X
9827	Power cord - Israel (250V, 6-16A)	X	X	X	X	X	X	X
9828	Power cord - Switzerland (250V, 10A)	X	X	X	X	X	X	X
9829	Power cord - South Africa/Pakistan (250V, 16A)	X	X	X	X	X	X	X
9830	Power cord - Italy (250V, 10A and 16A)	X	X	X	X	X	X	X
9831	Power cord - Argentina (250V, 10A)	X	X	X	X	X	X	X
9833	Power cord - Thailand (250V, 15A)	X	X	X	X	X	X	X
9834	Power cord - Uruguay (250V, 10A)	X	X	X	X	X	X	X
9900	Power cord (4M) - United States/Canada	X	X		X		X	
9901	Power cord (4M) - Belgium, Finland, France (250V, 16A)	X	X		X		X	
9902	Power cord (4M) - Denmark (250V, 10A)	X	X		X		X	
9903	Power cord (4M) - U.K. and Others (250V, 13A)	X	X		X		X	
9904	Power cord (4M) - Israel (250V, 6-16A)	X	X		X		X	
9905	Power cord (4M) - Switzerland (250V, 10A)	X	X		X		X	
9906	Power cord (4M) - India, Pakistan, S. Africa (250V, 16A)	X	X		X		X	
9907	Power cord (4M) - Italy (250V, 10A and 16A)	X	X		X		X	
9908	Power cord (4M) - Australia/New Zealand/Argentina (250V, 10A)	X	X		X		X	
9909	Power cord (4M) - Thailand (250V, 15A)	X	X		X		X	
9911	Power cord (4M) - All (Standard rack power cord)	X	X		X		X	

Table C-1 on page 908 lists the feature codes of the common power cords for Models 6M2, 651, 671, 681, 150, 170, 265, and 275.

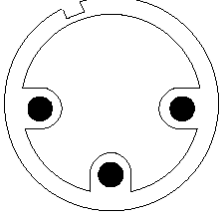
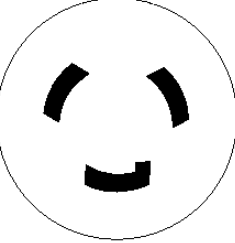
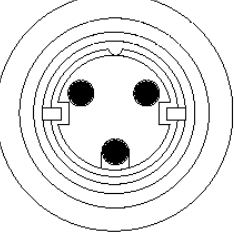
Table C-2 Power cord list by feature code

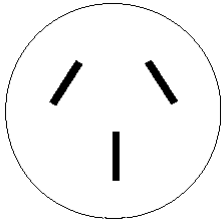
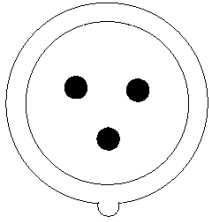
Feature code	Description	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
8677	Line Cord, WT, 50A/240V, 30A/480V, 8AWG, 14 ft., No Plug		X	X	X				
8678	Line Cord, 60A, 240V, 6AWG, 14 ft., US, Chicago IEC309 Plug			X	X				
8680	Line Cord, US, Canada, Japan, 480V, 30A, 10AWG, 14 ft., IEC309 Plug		X	X	X				
8681	Line Cord, US, Canada, Japan, 240V, 60A, 6AWG, 6 ft., Chicago IEC309 Plug			X	X				
8682	Line Cord, US, Chicago, 480V, 30A, 10AWG, 6 ft., Chicago, IEC309 Plug		X	X	X				
8683	Line Cord, US, Canada, Japan, 280V, 60A, 8AWG, 14 ft., IEC309 Plug, Single Phase			X	X				
8684	Line Cord, 8AWG, 6 ft., Chicago, IEC309 Plug, Single Phase, 60A, 200-280V			X	X				
8685	Line Cord, WT, 40A, 200-415V, 8AWG, 14 ft., No Plug, Single Phase			X	X				
8686	Line Cord, US, Canada, Japan, 200-240V, 6AWG, 14 ft., IEC309 100A Plug		X						
8687	Line Cord, US, Chicago, 200-240V, 6AWG, 6 ft., IEC309 100A Plug		X						
8688	Line Cord, US, Canada, Japan, 200-240V, 6AWG/Type W, 14 ft., IEC309 60A Plug		X						
8689	Line Cord, US, Chicago, 200-240V, 6AWG/Type W, 6 ft., IEC309 60A Plug		X						
9800	Power Cord - United States/Canada, L6-30P, 250V, single phase, 14 ft., 30A	X	X	X	X	X	X	X	X
9802	Power Cord - Brazil (125V, 15A)	X	X	X	X	X	X	X	X
9820	Power Cord - Belgium, Finland, France (250V, 16A)	X	X	X	X	X	X	X	X
9821	Power Cord - Denmark (250V, 10A)	X	X	X	X	X	X	X	X

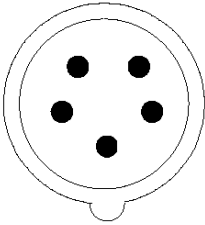
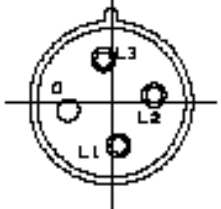
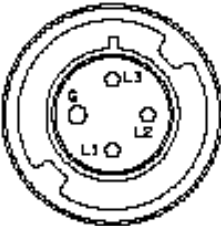
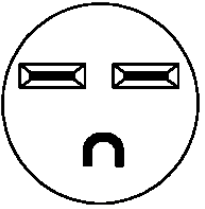
Feature code	Description	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
9825	Power Cord - United Kingdom (250V, 13A)	X	X	X	X	X	X	X	X
9827	Power Cord - Israel (250V, 6-16A)	X	X	X	X	X	X	X	X
9828	Power Cord - Switzerland (250V, 10A)	X	X	X	X	X	X	X	X
9829	Power Cord - South Africa/Pakistan (250V, 16A)	X	X	X	X	X	X	X	X
9830	Power Cord - Italy (250V, 10A and 16A)	X	X	X	X	X	X	X	X
9831	Power Cord - Argentina (250V, 10A)	X	X	X	X	X	X	X	X
9833	Power Cord - Thailand (250V, 15A)	X	X	X	X	X	X	X	X
9834	Power Cord - Uruguay (250V, 10A)	X	X	X	X	X	X	X	X
9900	Power Cord (4M) - United States/Canada								
9901	Power Cord (4M) - Belgium, Finland, France (250V, 16A)	X							
9902	Power Cord (4M) - Denmark (250V, 10A)	X							
9903	Power Cord (4M) - U.K. and Others (250V, 13A)	X							
9904	Power Cord (4M) - Israel (250V, 6-16A)	X							
9905	Power Cord (4M) - Switzerland (250V, 10A)	X							
9906	Power Cord (4M) - India, Pakistan, S. Africa (250V, 16A)	X							
9907	Power Cord (4M) - Italy (250V, 10A and 16A)	X							
9908	Power Cord (4M) - Australia/New Zealand/Argentina (250V, 10A)	X							

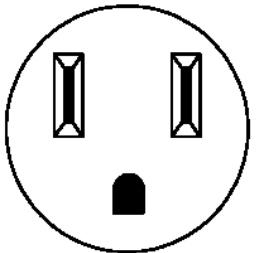
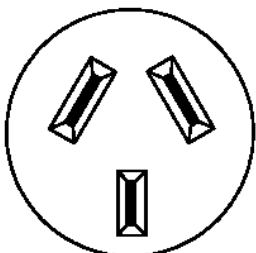
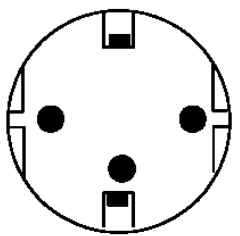
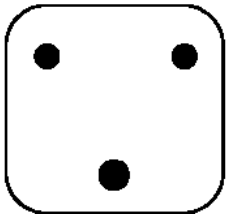
Table C-3 on page 911 provides a list of common power cords for main outlet connections, a picture of their connectors (not to scale), and country information.

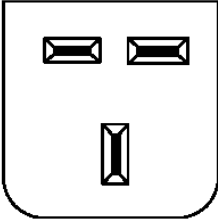


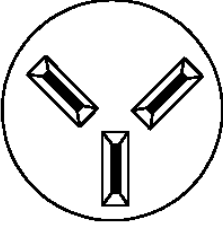
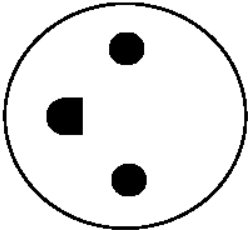
Table C-3 Power cord connector spotters reference

Picture	Feature code	Part number	Description
	9822	11F0106	For System Rack or Input/Output Rack, Wilco WP, type PDL, 250V, single phase, 14 ft., 30A Australia
	9826	11F0107	For System Rack or Input/Output Rack, Wilco WP, type PDL, 250V, single phase, 14 ft., 30A, right angle New Zealand
	9800, 9824	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.
	9986	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, Afghanistan, Argentina, Aruba, Bahamas, Bangladesh, Barbados, Belize, Bermuda, Bolivia, Bonaire, Brunei, Caicos Islands, Chile, Colombia, Costa Rica, Curaçao, Dominican Republic, Ecuador, El Salvador, Guatemala, Hong Kong S.A.R. of China, Indonesia, Jamaica, Japan, Malaysia, Mexico, Myanmar, Netherlands Antilles, Nicaragua, Panama, Paraguay, People's Republic of China, Peru, Philippines, St. Martin, Singapore, Sri Lanka, Suriname, Taiwan, Thailand, Trinidad/Tobago, Turks Island, Uruguay, Venezuela
	9987	46F4593	For System Rack or Input/Output Rack, 3 pin waterproof connector, type 40 plug, R&S 3750, 250V, single phase, 6 ft., 30A Chicago, Illinois U.S.
	9801	46F4594	For System Rack or Input/Output Rack, 3 pin waterproof connector, type 40 plug, R&S 3750, 250V, single phase, 14 ft., 30A U.S.A, Canada

Picture	Feature code	Part number	Description
	9835	87G6067	For System Rack or Input/Output Rack, type KP, right angle, 250V, single phase, 14 ft., 30A Korea
	9823	21H7693	For System Rack or Input/Output Rack, IEC 309, type 46 (2P+G), 250V, single phase, 14ft., 32A Abu Dubai, Albania, Algeria, Angola, Armenia, Bahrain, Belarus, Belgium, Benin, Bosnia, Botswana, Brazil, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Congo, Croatia, Cyprus, Djibouti, Dominica, Egypt, Equatorial Guinea, Ethiopia, France, French Guiana, Gabon, Gambia, Georgia, Germany, Ghana, Greece, Grenada, Grenadines, Guinea, Guinea-Bissau, Guyana, Haiti, Hungary, Iceland, India, Iran, Iraq, Ireland, Italy, Ivory Coast, Jordan, Kazakhstan, Kenya, Kirhizia, Kuwait, Lesotho, Lebanon, Liberia, Libya, Luxembourg, Macao S.A.R. of China, Macedonia, Malawi, Mali, Malta, Martinique, Mauritania, Moldova, Monaco, Morocco, Mozambique, Namibia, Nepal, Netherlands, New Caledonia, Niger, Nigeria, North Yemen, Oman, Pakistan, Poland, Principe, Qatar, Reunion, Romania, Rwanda, St. Lucia, St. Vincent, Sao Thome Island, Saudi Arabia, Senegal, Serbia, Sierra Leone, Slovakia, Slovenia, Somalia, South Yemen, Spain, Sudan, Swaziland, Syria, Tanzania, Togo, Tunisia, Uganda, Ukraine, United Arab Emirates, United Kingdom, Vietnam, Western Samoa, Zaire, Zambia, Zimbabwe

Picture	Feature code	Part number	Description
	6173, 9173	88G4763	For Input/Output Rack, IEC 309, type 46 (3P+N+G), 250V, two of three phase, 14 ft., 32A Austria, Czech Republic, Denmark, Estonia, Finland, Israel, Latvia, Liechtenstein, Lithuania, Netherlands, Norway, Pakistan, Portugal, Russia, South Africa, Sweden, Turkey
	6174, 9174	88G4764	For Input/Output Rack, IEC 309, type 46 (3P+N+G), 250V, three phase, 14 ft., 16A Switzerland
	8680	11P0916	Line Cord, US, Canada, Japan, 480V, 30A, 10AWG, 14 ft., IEC309 Plug
	8682	11P0914	Line Cord, US, Chicago, 480V, 30A, 10AWG, 6 ft., IEC309 Plug
	8683	44P1370	Line Cord, US, Canada, Japan, 200-280V, 6AWG, 14 ft., IEC309, 1 Phase, 60A Plug
	8678	11P0365	Line Cord, 6AWG, 14 ft., US, Chicago IEC309 Plug
	8681	11P0367	Line Cord, 6AWG, 6 ft., Chicago IEC309 Plug
	8684	44P1369	Line Cord, 8AWG, 6 ft., Chicago, IEC309 Plug, Single Phase, 60A, 200-280V
	8686	44P2290	Line Cord, US, Canada, Japan, 200-240V, 6AWG, 14 ft., IEC309 100A Plug
	8687	44P2289	Line Cord, US, Chicago, 200-240V, 6AWG/Type W, 6 ft., IEC309 100A Plug
	8688	44P2436	Line Cord, US, Canada, Japan, 200-240V, 6AWG, 14 ft., IEC309 60A Plug
	8689	44P2435	Line Cord, US, Chicago, 200-240V, 6AWG/Type W, 6 ft., IEC309 60A Plug
	9833	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	Feature code	Part number	Description
	9800 9986 9802	88G7648 62X1045 49P2103	Bahamas, Barbados, Bermuda, Bolivia, 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 Brazil
	9831, 9834	6952311 6952291	Argentina, Australia, New Zealand Paraguay, Colombia, Uruguay
	9820	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
	9829	14F0015	Bangladesh, Burma, Pakistan, South Africa, Sri Lanka

Picture	Feature code	Part number	Description
	9825	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
	9828	14F0051	Liechtenstein, Switzerland
	9830	14F0069	Chile, Ethiopia, Italy
	9827	14F0087	Israel
	9821	13F9997	Denmark



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

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 follows:

- ▶ SCSI-1: Maximum transfer rate of 5 MBps (one byte transfers)
- ▶ SCSI-2 Fast: Maximum transfer rate of 10 MBps (one byte transfers)
- ▶ SCSI-2 Fast/Wide: Maximum transfer rate of 20 MBps (two byte transfers)
- ▶ Ultra SCSI: Maximum transfer rate of 20 MBps (one byte transfers)
- ▶ Ultra SCSI Wide: Maximum transfer rate of 40 MBps (two byte transfers)

For this adapter:

- ▶ The maximum supported cable length for configurations without any SCSI-2 fast or ultra SCSI devices is 6 meters (approximately 20 feet).
- ▶ The maximum supported cable length for configurations that include SCSI-2 fast (but not ultra) is 3 meters (approximately 10 feet) with the exception of the 7027 High-Capacity Storage Drawer, which can be attached with up to 6 meters of cable.
- ▶ To ensure optimum signal quality for Ultra SCSI transfers, attachment of multiple Ultra SCSI devices is only recommended for devices mounted inside the system unit. This adapter has circuitry that can detect the presence of a cable on the external connector, and the default configuration limits the SCSI bus speed to SCSI-2 fast and wide (20 MBps).
- ▶ To ensure optimum signal quality for Ultra SCSI transfers, it is recommended that only Ultra SCSI devices be attached to a backplane that is driven by a PCI Single-Ended Ultra SCSI adapter.

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

Supported Ultra configurations for this adapter

The supported configurations for both internal and external devices connected to the PCI Single-Ended Ultra SCSI adapters are as follows:

- ▶ Internal Ultra devices running at Ultra speeds:
 - Up to six ultra devices attached to the internal port (dependent on internal configuration and cabling).
 - No external attachments are allowed.
- ▶ External Ultra devices running at Ultra speeds:
 - No internal attachments are allowed.
 - Up to two external Ultra wide (16-bit) devices can be attached to the external port. Maximum cable length must not exceed 3 meters.
 - Up to two external Ultra (8-bit) devices can be attached to the external port. Maximum cable length must not exceed 3 meters.
- ▶ Ultra or SCSI-2 Fast devices running at SCSI-2 Fast speeds:
 - Up to six devices can be attached to the internal port (dependant on internal system configuration and cabling).
 - External attachments of up to four independent physical enclosures are allowed, provided each physical enclosure presents only one load to the SCSI bus. The total bus length must not exceed 3 meters. Total bus length includes internal and external cable length.
- ▶ Multiple SCSI-2 Fast devices in external enclosures:
 - No internal attachments are allowed.
 - The maximum combined internal (to enclosure) and external cable length is 3 meters.
 - Loads on the cable must be 0.1 meters apart at a minimum (cable length between devices must be 0.1 meters apart at a minimum).
 - No mixing of bus widths (8-bit and 16-bit) unless the 68-pin to 50-pin interposer (P/N 92F2565 or equivalent) is used.

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)

Single-Ended adapter-to-first device cables

Table D-1 describes the cables for the Single-Ended Ultra SCSI Adapter for connection to the first device.

Table D-1 Single-Ended Ultra SCSI adapter-to-first device cables

Machine type	Feature code	Part number	Length (meters)	Cable description
Host System	2111	06H6037	1.0	Adapter-to-first device (where first device has two connectors), 8-bit narrow bus
Host System	2113	52G0174	1.5	Adapter-to-first device (where first device has one connector), 8-bit narrow bus
Host System	2115	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus

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

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

Device-to-device cables

Table D-2 describes the device-to-device cables for single-ended applications.

Table D-2 Device-to-device cables for single-ended installations

Machine type	Feature code	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

Machine type	Feature code	Part number	Length (meters)	Cable description
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.

System-to-system cables

Table D-3 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 D-3 16 Bit SCSI-II system-to-system cable

SE or DE device type	Feature code	Length (meters)	Cable description
External Devices	2424	0.6	16-bit SCSI-II system-to-system cable
External Devices	2425	2.5	16-bit SCSI-II system-to-system cable

Terminators for use with this adapter

Table D-4 describes the terminators for single-ended installations.

Table D-4 Terminators for single-ended installations

Machine type	Feature code	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

Cabling examples for the PCI Ultra SCSI adapter

Figure D-1 shows how to cable single-ended internal devices.

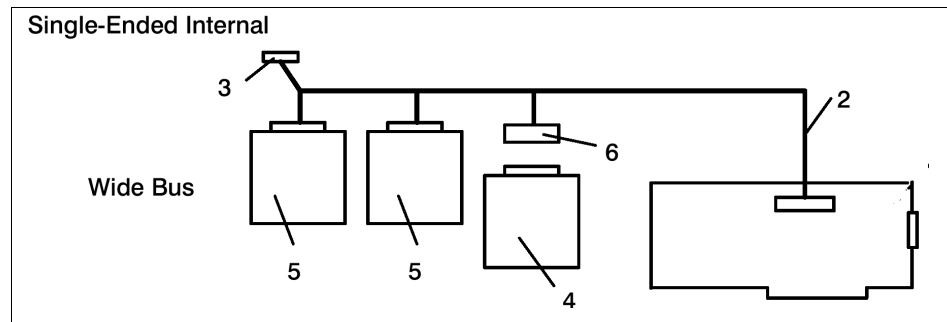


Figure D-1 Single-ended internal devices cabling

Table D-5 describes the internal system cables for use with this adapter.

Table D-5 Internal system cables

Item number	Description
1	Adapter
2	Internal SCSI cable and terminator assembly (wide)
3	System Specific Terminator (88G3977)
4	Narrow device (50-pin connector)

Item number	Description
5	Wide device (68-pin connector)
6	68-pin to 50-pin interposer FRU 92F0324 (assem. P/N 92F2565)

Figure D-2 shows how to cable external devices for both narrow and wide bus.




Figure D-2 External devices cabling for both narrow and wide bus

Table D-6 provides a list of available cables and terminators for single-ended adapters.

Table D-6 Cable and terminators for single-ended adapters

Item number	Part number	Description	Cable length (meters)
1		Adapter	N/A
2	06H6037	Adapter-to-dual-connector device (narrow 8-bit)	1.0
3	52G0174	Adapter-to-single-connector device (narrow 8-bit)	1.5
4	70G9857	Adapter-to-dual-connector device (wide 16-bit)	1.0

Item number	Part number	Description	Cable length (meters)
5	33F4607	Device-to-dual-connector device (narrow 8-bit)	0.7
6	31F4222	Device-to-single-connector device (narrow 8-bit)	0.66
7	52G9921	Device-to-dual-connector device (wide 16-bit)	0.3
8	52G4291	Device-to-dual-connector device (wide 16-bit)	0.6
9	52G4260	Terminator (8-bit)	
10	92F0432	Terminator (16-bit)	

Cabling the PCI Differential-Ended Ultra SCSI adapter

This section discusses the cabling for the PCI Differential Ultra SCSI Adapter (FC 6207).

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 (FC 6206)

Single-ended adapter-to-first device cables

Table D-7 on page 925 describes the cables for the Single-Ended Ultra SCSI Adapter for connection to the first device.

Table D-7 Single-Ended Ultra SCSI adapter-to-first device cables

Machine Type	Feature Code	Part Number	Length (meters)	Cable Description
Host System	2111	06H6037	1.0	Adapter-to-first device (where first device has two connectors), 8-bit narrow bus
Host System	2113	52G0174	1.5	Adapter-to-first device (where first device has one connector), 8-bit narrow bus
Host System	2115	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus

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

For 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 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.

SCSI Differential cable lengths using this adapter

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

Differential-ended adapter-to-first device cables

Table D-8 provides a list of adapter-to-first device cables.

Table D-8 Differential-ended Ultra SCSI adapter-to-first device cables

Machine type	Feature Code	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

Machine type	Feature Code	Part Number	Length (meters)	Cable description
Host System	2114	52G0173	0.94	16-bit Y-cable
Host System	2116	70G9857	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 C-10.

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

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

Device-to-device cables

Table D-9 provides a list of device-to-device cables.

Table D-9 Device-to-device cables for differential-ended installations

Machine type	Feature Code	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
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			

Machine type	Feature Code	Length (meters)	Part Number	Cable description
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, and 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.

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 D-10 provides a list of required terminators.

Table D-10 Terminator for differential-ended Installations

Machine type	Feature Code	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

Examples for the PCI Differential Ultra SCSI Adapter

Figure D-3 and Figure D-4 on page 929 on page 28 show how to cable the PCI Differential Ultra SCSI Adapter.

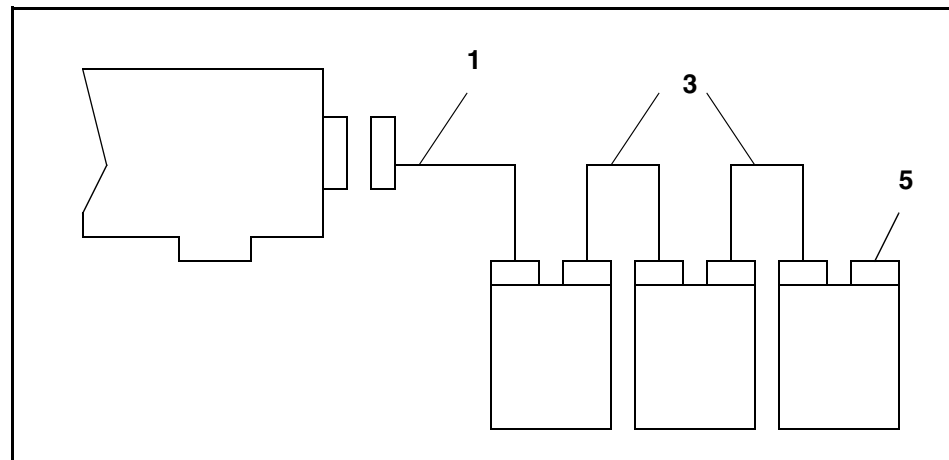


Figure D-3 Differential External Narrow Bus




Figure D-4 Differential External Wide Bus

Table D-11 provides a list of available cables and terminators for differential adapters.

Table D-11 Cables and terminators for the PCI Differential Ultra SCSI Adapter

Item Number	Part Number	Length (meters)	Description
1	06H6037	1.0	Adapter-to-dual-connector device (narrow 8-bit)
2	06H6036	1.0	Adapter-to-dual-connector device (wide 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)



Supported peripherals by device matrix

This appendix provides information about the supported peripheral by device for the following pSeries system models:

- ▶ 7026 Model B80
- ▶ 7028 Model 6C1 and 6E1
- ▶ 7028 Model 6C4 and 6E4
- ▶ 7029 Model 6C3 and 6E3
- ▶ 7038 Model 6M2
- ▶ 7039 Model 651
- ▶ 7040 Model 671
- ▶ 7040 Model 681
- ▶ 7043 Model 150
- ▶ 7044 Model 170
- ▶ 9112 Model 265
- ▶ 9114 Model 275

Table E-1 on page 932 and Table E-2 on page 938 summarize the external devices that are supported by current pSeries system models.

Table E-1 External devices for B80, 6C1, 6E1, 6C4, 6E4, 6C3, and 6E3

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
Switch							
2031-016 McDATA Departmental ES-3032 Fabric Switch		X	X	X	X	X	X
2031-032 McDATA Departmental ES-3032 Fabric Switch		X	X	X	X	X	X
2031-216 McData ES-3216 Fabric Switch				X	X	X	X
2031-224 McData Switch						X	
2031-232 McData ES-3232 Fabric Switch				X	X	X	X
2031-L00 McDATA Model L00 ES-1000 Loop Switch		X	X	X	X	X	X
2032-001 MCDATA ED-5000 Enterprise Fibre Channel	X	X	X	X	X	X	X
2032-064 MCDATA Enterprise Fibre Channel Model 064		X	X	X	X	X	X
2032-140 MCDATA Enterprise Fibre Channel Model 140						X	
2042-001 Inrange 64 port Fibre Channel Director		X	X	X	X	X	X
2042-128 Inrange 128 port Fibre Channel Director		X	X	X	X	X	X
2042-256 Inrange 256 port Fibre Channel Director				X	X	X	X
2062-D01 Cisco MDS-9216 Multilayer Fabric Switch						X	
2062-D07 Cisco MDS-9509 Multilayer Director Switch						X	
2062-T07 Cisco MDS-9509 Multilayer Director Switch (Vdc)						X	
2109-F16 IBM TotalStorage SAN Switch Model F16		X	X	X	X	X	X
2109-F32 IBM TotalStorage SAN Switch Model F32						X	
2109-M12 IBM TotalStorage SAN Switch Model M12				X	X	X	X
2109-S08 IBM SAN Fibre Channel Switch Model S08	X	X	X	X	X	X	X
2109-S16 IBM SAN Fibre Channel Switch Model S16	X	X	X	X	X	X	X
3534-F08 IBM TotalStorage SAN Switch Model F08						X	
Disk drives and subsystems							
2102-D00 Expandable Storage Unit	X	X	X	X	X	X	X
2102-F10 Fibre Channel RAID Storage Server	X	X	X	X	X	X	X
2104-DL1 Expandable Storage Plus	X	X	X	X	X	X	X
2104-DU3 Expandable Storage Plus Model DU3	X	X	X	X	X	X	X
2104-TL1 Expandable Storage Plus Model TL1	X	X	X	X	X	X	X
2104-TU3 Expandable Storage Plus Model TU3	X	X	X	X	X	X	X
2105-800 TotalStorage Enterprise Storage Server				X	X	X	X
2105-E10 Enterprise Storage Server Model E10	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
2105-E20 Enterprise Storage Server Model E20	X	X	X	X	X	X	X
2105-F10 Enterprise Storage Server Model F10	X	X	X	X	X	X	X
2105-F20 Enterprise Storage Server Model F20	X	X	X	X	X	X	X
3542-2RU Fast T200 Storage Server Model 2RU				X	X	X	X
3542-2RX Fast T200 Storage Server Model 2RX				X	X	X	X
3552-1RU TotalStorage Fast T500 Storage Server				X	X	X	X
3552-1RX TotalStorage Fast T200 Storage Server				X	X	X	X
3560-1RU TotalStorage Fast EXP500 Storage Expansion				X	X	X	X
3560-1RX TotalStorage Fast EXP500 Storage Expansion				X	X	X	X
7131-105 Multi-Storage Tower Model 105	X	X	X	X	X	X	X
7131-405 Multi-Storage Tower Model 405	X	X		X	X	X	X
7133-010 SSA Disk Subsystem Model 010	X	X		X	X	X	X
7133-020 SSA Disk Subsystem Model 020	X	X		X	X	X	X
7133-500 SSA Disk Subsystem Model 500	X	X		X	X	X	X
7133-600 SSA Disk Subsystem Model 600	X	X		X	X	X	X
7133-D40 SSA Disk Subsystem Model D40	X	X	X	X	X	X	X
7133-T40 SSA Disk Subsystem Model T40	X	X	X	X	X	X	X
7137-412 Disk Array Subsystem Model 412	X	X	X	X	X	X	X
7137-413 Disk Array Subsystem Model 413	X	X	X	X	X	X	X
7137-414 Disk Array Subsystem Model 414	X	X	X	X	X	X	X
7137-415 Disk Array Subsystem Model 415	X	X	X	X	X	X	X
7203-001 External Portable Disk Drive Model 001	X	X	X	X	X	X	X
7204-010 External 1.0 GB Disk Drive Model 010	X	X	X	X	X	X	X
7204-112 External 1.1 GB F/W Disk Drive Model 112	X	X	X	X	X	X	X
7204-113 External Disk Drive Model 113	X	X	X	X	X	X	X
7204-114 External Disk Drive Model 114	X	X	X	X	X	X	X
7204-118 External 18.21 GB Disk Drive Model 118	X	X	X	X	X	X	X
7204-139 External Disk Drive Model 139	X	X	X	X	X	X	X
7204-215 External 2.0 GB Differential Disk Drive Model 112	X	X	X	X	X	X	X
7204-315 External Disk Drive Model 315	X	X	X	X	X	X	X
7204-317 External Disk Drive Model 317	X	X	X	X	X	X	X
7204-325 External Disk Drive Model 325	X	X	X	X	X	X	X
7204-339 External Disk Drive Model 339	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
7204-402 External Disk Drive Model 402	X	X	X	X	X	X	X
7204-404 External Disk Drive Model 404	X	X	X	X	X	X	X
7204-409 External 9.1 GB Disk Drive Model 409	X	X	X	X	X	X	X
7204-418 External 18 GB Differential Disk Drive Model 418	X	X	X	X	X	X	X
7204-419 External 18.2 GB Disk Drive Model 419	X	X	X	X	X	X	X
7204-518 External 18.3 GB Disk Drive Model 518	X	X	X	X	X	X	X
7204-536 External 36.7 GB Disk Drive Model 536	X	X	X	X	X	X	X
Graphics related							
6093-011 CursorPad Tablet	X	X	X	X	X	X	X
6093-012 Large Tablet	X	X	X	X	X	X	X
6093-021 Tablet Model 021	X	X	X	X	X	X	X
6094-010 Dials	X	X	X	X	X	X	X
6094-020 LPF Keyboard	X	X	X	X	X	X	X
6094-030 Spaceball	X	X	X	X	X	X	X
6094-031 Spaceball	X	X	X	X	X	X	X
6094-040 Spaceball	X	X	X	X	X	X	X
6094-600 Magellan Enhanced Input Device	X	X	X	X	X	X	X
7250-002 POWER GXT1000 Graphics Accelerator	X	X	X	X	X	X	X
Displays							
6546-00E G52 Display	X	X	X	X	X	X	X
6546-00N G52 Display	X	X	X	X	X	X	X
6546-01E G52 Display	X	X	X	X	X	X	X
6546-01S G52 Display	X	X	X	X	X	X	X
6546-0AE G54 Display	X	X	X	X	X	X	X
6546-0AN G54 Display	X	X	X	X	X	X	X
6546-0BE G54 Display	X	X	X	X	X	X	X
6546-0BN G54 Display	X	X	X	X	X	X	X
6546-0BS G54 Display	X	X	X	X	X	X	X
6546-21N G52 Display	X	X	X	X	X	X	X
6546-2BN G54 Display	X	X	X	X	X	X	X
6546-31N G52 Display	X	X	X	X	X	X	X
6546-4AN G54 Display	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
6546-6BN G54 Display	X	X	X	X	X	X	X
6555-763 P70 17-in. Color Display	X	X	X	X	X	X	X
6555-764 P70 17-in. Color Display	X	X	X	X	X	X	X
6555-773 P200 UV-N Color Display	X	X	X	X	X	X	X
6555-774 P200 UV-N Color Display	X	X	X	X	X	X	X
6555-803 P201 UV-N Color Display	X	X	X	X	X	X	X
7515-C01 Hardware Management Console				X		X	
7316-TF1 Rack Mounted Flat Panel Display		X	X	X	X	X	X
7316-TF2 Rack Mounted Flat Panel Display				X	X	X	
9513-AG2 T55A 15-in. TFT LCD Color Monitor	X	X	X	X	X	X	X
9513-AW1 T55A 15-in. TFT LCD Color Monitor	X	X	X	X	X	X	X
9516-B03 16.1-in. TFT LCD Display	X	X	X	X	X	X	X
9516-B04 16.1-in. TFT LCD Display	X	X	X	X	X	X	X
Storage network							
1722-60U FAStT600 Storage Network						X	
1740-1RU TotalStorage Fast EXP700 Expansion Unit						X	
1742-1RU TotalStorage Fast EXP700 Storage Server						X	
1742-90U FAStT900 Storage Network						X	
2108-G07 SAN Data Gateway Model G07	X	X	X	X	X	X	X
2108-R03 SAN Data Gateway Router	X	X	X	X	X	X	X
Optical drives and libraries							
3995-C60 Optical Library Dataserver	X	X	X	X	X	X	X
3995-C62 Optical Library Dataserver	X	X	X	X	X	X	X
3995-C64 Optical Library Dataserver	X	X	X	X	X	X	X
3995-C66 Optical Library Dataserver	X	X	X	X	X	X	X
3995-C68 Optical Library Dataserver	X	X	X	X	X	X	X
7209-002 Optical Disk Drive Model 002	X	X	X	X	X	X	X
7209-003 Multifunction Optical Disk Drive Model 003	X	X	X	X		X	
7210-005 CD-ROM Drive Model 005	X	X	X	X	X	X	X
7210-010 CD-ROM Drive Model 010	X	X	X	X	X	X	X
7210-015 CD-ROM Drive Model 015	X	X	X	X	X	X	X
7210-020 CD-ROM Drive Model 020	X	X	X	X	X	X	X
7210-025 DVD-RAM Drive Model 025	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
7212-102 Storage Device Enclosure (SAM BASS)		X	X	X	X	X	X
Tape drives and libraries							
3490-F00 Magnetic Tape Subsystem Model F00	X	X	X	X	X	X	X
3490-F01 Magnetic Tape Subsystem Model F01	X	X	X	X	X	X	X
3490-F11 Magnetic Tape Subsystem Model F11	X	X	X	X	X	X	X
3494-B10 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X
3494-B18 Magstar Virtual Tape Server		X	X	X	X	X	X
3494-B20 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X
3494-D12 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X
3494-L12 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X
3570-B00 Magstar MP Tape Subsystem Model B00	X	X	X	X	X	X	X
3570-B01 Magstar MP Tape Subsystem Model B01	X	X	X	X	X	X	X
3570-B02 Magstar MP Tape Subsystem Model B02	X	X	X	X	X	X	X
3570-B11 Magstar MP Tape Subsystem Model B11	X	X	X	X	X	X	X
3570-B12 Magstar MP Tape Subsystem Model B12	X	X	X	X	X	X	X
3570-B1A Magstar MP Tape Drive Model B1A	X	X	X	X	X	X	X
3570-C00 Magstar MP Tape Drive Model C00	X	X	X	X	X	X	X
3570-C01 Magstar MP Tape Drive Model C01	X	X	X	X	X	X	X
3570-C02 Magstar MP Tape Drive Model C02	X	X	X	X	X	X	X
3570-C11 Magstar MP Tape Drive Model C11	X	X	X	X	X	X	X
3570-C12 Magstar MP Tape Drive Model C12	X	X	X	X	X	X	X
3575-L06 Magstar MP Tape Library Dataserver Model L06	X	X	X	X	X	X	X
3575-L12 Magstar MP Tape Library Dataserver Model L12	X	X	X	X	X	X	X
3575-L18 Magstar MP Tape Library Dataserver Model L18	X	X	X	X	X	X	X
3575-L24 Magstar MP Tape Library Dataserver Model L24	X	X	X	X	X	X	X
3575-L32 Magstar MP Tape Library Dataserver Model L32	X	X	X	X	X	X	X
3580-H11 Linear Tape Open (LTO) Bridge Box (HVD)	X	X	X	X	X	X	X
3580-H13 Ultrium LTO Tape Drive				X	X	X	X
3580-H23 TotalStorage Ultrium 2 Tape Drive	X	X	X	X	X	X	X
3580-L11 Linear Tape Open (LTO) Bridge Box (LVD)	X	X	X	X	X	X	X
3580-L13 Ultrium LTO Autoloader				X	X	X	X
3580-L23 TotalStorage Ultrium 2 Tape Drive	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
3581-H13 Ultrium LTO Tape Drive				X	X	X	X
3581-H17 Linear Tape Open (LTO) Autoloader (HVD)	X	X	X	X	X	X	X
3581-L13 Ultrium LTO Tape Autoloader				X	X	X	X
3581-L17 Linear Tape Open (LTO) Autoloader (LVD)	X	X	X	X	X	X	X
3583-L18 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X
3583-L36 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X
3583-L72 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X
3584-D32 Ultrium Ultra-Scalable Tape Library - Base Frame	X	X	X	X	X	X	X
3584-D42 Ultra-Scalable Tape Library	X	X	X	X	X	X	X
3584-L32 Ultrium Ultra-Scalable Tape Library - Expansion	X	X	X	X	X	X	X
3590-B11 Magstar Tape Subsystem Model B11	X	X	X	X	X	X	X
3590-B1A Magstar Tape Subsystem Model B1A	X	X	X	X	X	X	X
3590-C12 Magstar Silo Compatible Model C12	X	X	X	X	X	X	X
3590-E11 Magstar Tape Subsystem Model E11	X	X	X	X	X	X	X
3590-E1A Magstar Tape Subsystem Model E1A	X	X	X	X	X	X	X
3590-H11 TotalStorage Enterprise Tape Drive Model H11	X	X	X	X	X	X	X
3590-H1A TotalStorage Enterprise Tape Drive Model H1A	X	X	X	X	X	X	X
7205-311 35 GB External DLT Tape Drive Model 311	X	X	X	X	X	X	X
7205-440 40 GB External DLT Tape Drive Model 440	X		X	X	X		X
7205-550 External DLT Tape Drive Model 550						X	X
7206-005 External 4 mm Tape Drive Model 005	X	X	X	X	X	X	X
7206-110 12 GB External 4 mm DDS-3 Tape Drive Model 110	X	X	X	X	X	X	X
7206-220 20 GB External 4 mm DDS-4 Tape Drive Mod. 220	X	X	X	X	X	X	X
7206-VX2 External Tape Drive Model VX2	X	X	X	X	X	X	X
7207-011 1/4-in. External Cartridge Tape Drive Model 011	X	X	X	X		X	
7207-012 1/4-in. External Cartridge Tape Drive Model 012	X	X	X	X	X	X	X
7207-122 4 GB 1/4-in. External SLR5 QIC Tape Drive	X	X	X	X	X	X	X
7207-315 13 GB 1/4-in. External Cartridge Tape Drive	X	X	X	X	X	X	X
7208-011 External 8 mm Tape Drive	X	X	X	X	X	X	X
7208-341 20 Gb External 8 mm Tape	X	X	X	X	X	X	X
7208-345 60 GB External 8 mm Tape	X	X	X	X	X	X	X
7331-205 External 8 mm Tape Library	X	X	X	X	X	X	X

External supported devices	Machine models						
	B 8 0	6 C 1	6 E 1	6 C 4	6 E 4	6 C 3	6 E 3
7331-305 External 8 mm Tape Library	X	X	X	X	X	X	X
7332-005 4 mm DDS-2 Tape Autoloader	X	X	X	X	X	X	X
7332-110 4 mm DDS-3 Tape Autoloader	X	X	X	X	X	X	X
7332-220 4 mm DDS-3 Tape Autoloader	X	X	X	X	X	X	X
7334-410 8 mm Tape Library	X	X	X	X	X	X	X
7336-205 4 mm Tape Library	X	X	X	X	X	X	X
7337-305 DLT Library	X	X	X	X	X	X	X
7337-306 DLT Library	X	X	X	X	X	X	X
7337-360 DLT Library	X	X	X	X	X	X	X
9348-012 Magnetic Tape Unit Model 012	X	X	X	X	X	X	X
Expansion cabinets							
2032-C36 McDATA Cabinet for Directors and Fabric Switch		X	X	X	X	X	X
2042-C40 Inrange 40U Cabinet		X	X	X	X	X	X
2109-C36 TotalStorage San Cabinet				X	X	X	X
7014-S00 1.6 Meter System Rack - 32 U	X	X	X	X	X	X	X
7014-T00 1.8 Meter System Rack - 36 U	X	X	X	X	X	X	X
7014-T42 2.0 Meter System Rack - 42 U	X	X	X	X	X	X	X
7015-R00 1.6 Meter System Rack - 32 U							
7311-D10 I/O Drawer							
7311-D20 I/O Drawer Rack -Mounted Expansion Drawer				X		X	

Table E-2 External devices for 6M2, 651, 671, 681, 150, 170, 265, and 275

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
Switch								
2031-016 McDATA Departmental ES-3032 Fabric Switch	X	X	X	X		X	X	X
2031-032 McDATA Departmental ES-3032 Fabric Switch	X	X	X	X			X	X
2031-216 McData ES-3216 Fabric Switch	X	X						
2031-224 McData ES-3224 Fabric Switch	X							
2031-232 McData ES-3232 Fabric Switch	X	X						
2031-L00 McDATA Model L00 ES-1000 Loop Switch	X	X	X	X		X	X	X
2032-001 MCDATA ED-5000 Enterprise Fibre Channel	X	X	X	X		X	X	X
2032-064 MCDATA Enterprise Fibre Channel Model 064	X	X	X	X		X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
2032-140 MCDATA Enterprise Fibre Channel Model 140	X							
2042-001 Inrange 64-port Fibre Channel Director	X	X	X	X		X	X	X
2042-128 Inrange 128-port Fibre Channel Director	X	X	X	X		X	X	X
2042-256 Inrange 256-port Fibre Channel Director	X	X						
2109-F16 IBM TotalStorage SAN Switch Model F16	X	X	X	X			X	X
2109-F32 IBM TotalStorage SAN Switch Model F32	X	X						
2109-M12 IBM TotalStorage SAN Switch Model M12	X	X	X					
2109-S08 IBM SAN Fibre Channel Switch Model S08	X	X	X	X		X	X	X
2109-S16 IBM SAN Fibre Channel Switch Model S16	X	X	X	X		X	X	X
3534-F08 TotalStorage SAN Model F08	X	X						
Disk drives and subsystems								
2101-100 Seascope® Solution rack	X	X	X	X				
2102-D00 Expandable Storage Unit	X	X	X	X				
2102-F10 Fibre Channel RAID Storage Server	X	X	X	X				
2103-H07 Fibre Channel Storage Hub	X							
2104-DL1 Expandable Storage Plus	X	X	X	X	X	X	X	X
2104-DU3 Expandable Storage Plus Model DU3	X	X	X	X	X	X	X	X
2104-TL1 Expandable Storage Plus Model TL1	X	X	X	X	X	X	X	X
2104-TU3 Expandable Storage Plus Model TU3	X	X	X	X	X	X	X	X
2105-800 Totals tor age Enterprise Storage Server Model E10	X							
2105-E10 Enterprise Storage Server Model E10	X	X	X	X		X		
2105-E20 Enterprise Storage Server Model E20	X	X	X	X		X		
2105-F10 Enterprise Storage Server Model F10	X	X	X	X		X		
2105-F20 Enterprise Storage Server Model F20	X	X	X	X		X		
3542-2RU Fast T200 Storage Server Model 2RU	X							
3542-2RX Fast T200 Storage Server Model 2RX	X							
3552-1RU TotalStorage Fast T500 Storage Server	X						X	X
3552-1RX TotalStorage Fast T200 Storage Server	X						X	X
3560-1RU TotalStorage Fast EXP500 Storage Expansion	X							
3560-1RX TotalStorage Fast EXP500 Storage Expansion	X							
7027-HSC High Capacity Drawer Model HSC		X	X	X				

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
7027-HSD High Capacity Drawer Model HSD		X	X	X				
7131-105 Multi-Storage Tower Model 105		X	X	X	X	X	X	X
7131-405 Multi-Storage Tower Model 405		X	X	X	X	X	X	X
7133-010 SSA Disk Subsystem Model 010	X	X	X	X	X	X	X	X
7133-020 SSA Disk Subsystem Model 020	X	X	X	X	X	X	X	X
7133-500 SSA Disk Subsystem Model 500	X	X	X	X	X	X	X	X
7133-600 SSA Disk Subsystem Model 600	X	X	X	X	X	X	X	X
7133-D40 SSA Disk Subsystem Model D40	X	X	X	X	X	X	X	X
7133-T40 SSA Disk Subsystem Model T40	X	X	X	X	X	X	X	X
7134-010 High Density SCSI Disk Subsystem			X	X				
7135-210 RAIDiant Array		X	X	X				
7137-412 Disk Array Subsystem Model 412		X	X	X	X	X	X	X
7137-413 Disk Array Subsystem Model 413		X	X	X	X	X	X	X
7137-414 Disk Array Subsystem Model 414		X	X	X	X	X	X	X
7137-415 Disk Array Subsystem Model 415		X	X	X	X	X	X	X
7137-512 Disk Array Subsystem Model 512		X	X	X				
7137-513 Disk Array Subsystem Model 513		X	X	X				
7137-514 Disk Array Subsystem Model 514		X	X	X				
7137-515 Disk Array Subsystem Model 515		X	X	X				
7203-001 External Portable Disk Drive Model 001		X	X	X	X	X	X	X
7204-001 External 1.0 GB Disk Drive		X	X	X				
7204-010 External 1.0 GB Disk Drive Model 010		X	X	X	X	X	X	X
7204-112 External 1.1 GB F/W Disk Drive Model 112		X	X	X	X	X	X	X
7204-113 External Disk Drive Model 113		X	X	X	X	X	X	X
7204-114 External Disk Drive Model 114		X	X	X	X	X	X	X
7204-118 External 18.21 GB Disk Drive Model 118		X	X	X	X	X	X	X
7204-139 External Disk Drive Model 139		X	X	X	X	X	X	X
7204-215 External 2.0 GB Differential Disk Drive Model 112		X	X	X	X	X	X	X
7204-315 External Disk Drive Model 315		X	X	X	X	X	X	X
7204-317 External Disk Drive Model 317		X	X	X	X	X	X	X
7204-320 External Disk Drive Model 320		X	X	X			X	X
7204-325 External Disk Drive Model 325		X	X	X	X	X	X	X
7204-339 External Disk Drive Model 339		X	X	X	X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
7204-402 External Disk Drive Model 402		X	X	X	X	X	X	X
7204-404 External Disk Drive Model 404		X	X	X	X	X	X	X
7204-409 External 9.1 GB Disk Drive Model 409	X	X	X	X	X	X	X	X
7204-418 External 18 GB Differential Disk Drive Model 418	X	X	X	X	X	X	X	X
7204-419 External 18.2 GB Disk Drive Model 419	X	X	X	X	X	X	X	X
7204-518 External 18.3 GB Disk Drive Model 518	X	X	X	X	X	X	X	X
7204-536 External 36.7 GB Disk Drive Model 536	X	X	X	X	X	X	X	X
Graphics related								
6093-011 CursorPad Tablet					X	X	X	X
6093-012 Large Tablet					X	X	X	X
6093-021 Tablet Model 021					X	X	X	X
6094-010 Dials					X	X	X	X
6094-020 LPF Keyboard					X	X	X	X
6094-030 Spaceball					X	X	X	X
6094-031 Spaceball					X	X	X	X
6094-040 Spaceball					X	X	X	X
6094-051 Spaceball					X		X	X
6094-600 Magellan Enhanced Input Device					X	X	X	X
7250-002 POWER GXT1000 Graphics Accelerator					X	X	X	X
Displays								
6546-00E G52 Display	X	X	X	X	X	X	X	X
6546-00N G52 Display	X	X	X	X	X	X	X	X
6546-01E G52 Display	X	X	X	X	X	X	X	X
6546-01S G52 Display	X	X	X	X	X	X	X	X
6546-0AE G54 Display	X	X	X	X	X	X	X	X
6546-0AN G54 Display	X	X	X	X	X	X	X	X
6546-0BE G54 Display	X	X	X	X	X	X	X	X
6546-0BN G54 Display	X	X	X	X	X	X	X	X
6546-0BS G54 Display	X	X	X	X	X	X	X	X
6546-21N G52 Display	X	X	X	X	X	X	X	X
6546-2BN G54 Display	X	X	X	X	X	X	X	X
6546-31N G52 Display	X	X	X	X	X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
6546-4AN G54 Display	X	X	X	X	X	X	X	X
6546-6BN G54 Display	X	X	X	X	X	X	X	X
6555-763 P70 17-in. Color Display					X	X	X	X
6555-764 P70 17-in. Color Display					X	X	X	X
6555-773 P200 UV-N Color Display					X	X	X	X
6555-774 P200 UV-N Color Display					X	X	X	X
6555-803 P201 UV-N Color Display					X	X	X	X
7315-C01 Hardware Management Console	X		X	X				
7316-TF1 Rack Mounted Flat Panel Display	X		X	X				
7316-TF2 Rack Mounted Flat Panel Display	X		X	X				
9513-AG2 T55A 15-in. TFT LCD Color Monitor					X	X	X	X
9513-AW1 T55A 15-in. TFT LCD Color Monitor					X	X	X	X
9516-B03 16.1-in. TFT LCD Display					X	X	X	X
9516-B04 16.1-in. TFT LCD Display					X	X	X	X
9519-AG1 18-in. T85A LCD Display - Black					X			
9519-AW1 18-in. T85A LCD Display - White					X			
Storage Network								
1740-1RU TotalStorage Fast EXP700 Expansion Unit	X							
1742-1RU TotalStorage Fast EXP700 Storage Server	X							
2108-G07 SAN Data Gateway Model G07		X	X	X	X	X		
2108-R03 SAN Data Gateway Router		X	X	X	X	X	X	X
Optical drives and libraries								
3995-063 Optical Library Dataserver (two Drives)	X		X	X				
3995-163 Optical Library Dataserver (four Drives)	X		X	X				
3995-A63 Optical Library Dataserver (one Drive)	X		X	X				
3995-C60 Optical Library Dataserver	X	X	X	X	X	X	X	X
3995-C62 Optical Library Dataserver	X	X	X	X	X	X	X	X
3995-C64 Optical Library Dataserver	X	X	X	X	X	X	X	X
3995-C66 Optical Library Dataserver	X	X	X	X	X	X	X	X
3995-C68 Optical Library Dataserver	X	X	X	X	X	X	X	X
7209-002 Optical Disk Drive Model 002						X	X	X
7209-003 Multifunction Optical Disk Drive Model 003		X	X	X	X	X	X	X
7210-005 CD-ROM Drive Model 005					X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
7210-010 CD-ROM Drive Model 010					X	X	X	X
7210-015 CD-ROM Drive Model 015					X	X	X	X
7210-020 CD-ROM Drive Model 020					X	X	X	X
7210-025 DVD-RAM Drive Model 025	X	X	X	X	X	X	X	X
7212-102 Storage Device Enclosure (SAM BASS)	X						X	X
Tape drives and libraries								
3490-C11 Magnetic Tape Subsystem Model C11		X	X	X				
3490-C1A Magnetic Tape Subsystem Model C1A		X	X	X				
3490-C22 Magnetic Tape Subsystem Model C22		X	X	X				
3490-C2A Magnetic Tape Subsystem Model C2A		X	X	X				
3490-E01 Magnetic Tape Subsystem Model E01		X	X	X				
3490-E11 Magnetic Tape Subsystem Model E11		X	X	X				
3490-F00 Magnetic Tape Subsystem Model F00	X	X	X	X	X	X	X	X
3490-F01 Magnetic Tape Subsystem Model F01	X	X	X	X	X	X	X	X
3490-F11 Magnetic Tape Subsystem Model F11	X	X	X	X				
3490-F1A Magnetic Tape Subsystem Model F1A	X	X	X	X				
3490-FC0 Tape Rack Model FC0	X	X	X	X				
3494-B10 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X	X
3494-B18 Magstar Virtual Tape Server	X	X	X	X		X	X	X
3494-B20 TotalStorage Virtual Tape Server	X	X	X	X	X	X	X	X
3494-D10 TotalStorage Virtual Tape Server		X	X	X				
3494-D12 TotalStorage Virtual Tape Server	X	X	X	X				
3494-D14 TotalStorage Virtual Tape Server	X	X	X	X				
3494-HA1 Magstar High Availability Unit	X	X	X	X				
3494-L10 Magstar Tape Library Model L10	X	X	X	X				
3494-L12 Magstar Tape Library Model L12	X	X	X	X		X	X	X
3494-L14 Magstar Tape Library Model L14	X	X	X	X				
3494-S10 Magstar Tape Library Model S10	X	X	X	X				
3570-B00 Magstar MP Tape Subsystem Model B00		X	X	X	X	X	X	X
3570-B01 Magstar MP Tape Subsystem Model B01		X	X	X	X	X	X	X
3570-B02 Magstar MP Tape Subsystem Model B02		X	X	X	X	X	X	X
3570-B11 Magstar MP Tape Subsystem Model B11		X	X	X	X	X	X	X
3570-B12 Magstar MP Tape Subsystem Model B12		X	X	X	X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
3570-B1A Magstar MP Tape Drive Model B1A		X	X	X	X	X	X	X
3570-C00 Magstar MP Tape Drive Model C00		X	X	X	X	X	X	X
3570-C01 Magstar MP Tape Drive Model C01	X	X	X	X	X	X	X	X
3570-C02 Magstar MP Tape Drive Model C02	X	X	X	X	X	X	X	X
3570-C11 Magstar MP Tape Drive Model C11	X	X	X	X	X	X	X	X
3570-C12 Magstar MP Tape Drive Model C12	X	X	X	X	X	X	X	X
3575-L06 Magstar MP Tape Library Dataserver Model L06	X	X	X	X	X	X	X	X
3575-L12 Magstar MP Tape Library Dataserver Model L12	X	X	X	X	X	X	X	X
3575-L18 Magstar MP Tape Library Dataserver Model L18	X	X	X	X	X	X	X	X
3575-L24 Magstar MP Tape Library Dataserver Model L24	X	X	X	X	X	X	X	X
3575-L32 Magstar MP Tape Library Dataserver Model L32	X	X	X	X	X	X	X	X
3580-H11 Linear Tape Open (LTO) Bridge Box (HVD)	X	X	X	X	X	X	X	X
3580-H13 Ultrium LTO Tape Drive	X	X						
3580-H23 TotalStorage Ultrium 2 Tape Drive	X	X	X	X	X	X	X	X
3580-L11 Linear Tape Open (LTO) Bridge Box (LVD)	X	X	X	X	X	X	X	X
3580-L13 Ultrium LTO Autoloader	X	X						
3580-L23 TotalStorage Ultrium 2 Tape Drive	X	X	X	X	X	X	X	X
3581-H13 Ultrium LTO Tape Drive	X	X						
3581-H17 Linear Tape Open (LTO) Autoloader (HVD)	X	X	X	X	X	X	X	X
3581-L13 Ultrium LTO Tape Autoloader	X	X						
3581-L17 Linear Tape Open (LTO) Autoloader (LVD)	X	X	X	X	X	X	X	X
3583-L18 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X	X
3583-L36 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X	X
3583-L72 Linear Tape Open (LTO) Tape Library - Scalable	X	X	X	X	X	X	X	X
3584-D32 Ultrium Ultra-Scalable Tape Library - Base Frame	X	X	X	X	X	X	X	X
3584-D42 Ultra-Scalable Tape Library	X	X	X	X	X	X	X	X
3584-L32 Ultrium Ultra-Scalable Tape Library - Expansion	X	X	X	X	X	X	X	X
3590-B11 Magstar Tape Subsystem Model B11	X	X	X	X	X	X	X	X
3590-B1A Magstar Tape Subsystem Model B1A	X	X	X	X	X	X	X	X
3590-C12 Magstar Silo Compatible Model C12	X	X	X	X	X	X	X	X
3590-E11 Magstar Tape Subsystem Model E11	X	X	X	X	X	X	X	X
3590-E1A Magstar Tape Subsystem Model E1A	X	X	X	X	X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
3590-H11 TotalStorage Enterprise Tape Drive Model H11	X	X	X	X	X	X	X	X
3590-H1A TotalStorage Enterprise Tape Drive Model H1A	X	X	X	X	X	X	X	X
7205-311 35 GB External DLT Tape Drive Model 311	X	X	X	X	X	X	X	X
7205-440 40 GB External DLT Tape Drive Model 440	X	X	X	X	X	X	X	X
7206-005 External 4 mm Tape Drive Model 005		X	X	X	X	X	X	X
7206-110 12 GB External 4 mm DDS-3 Tape Drive Model 110	X	X	X	X	X	X	X	X
7206-220 20 GB External 4 mm DDS-4 Tape Drive Mod. 220	X	X	X	X	X	X	X	X
7206-VX2 External Tape Drive Model VX2	X	X	X	X	X	X	X	X
7207-011 1/4-in. External Cartridge Tape Drive Model 011		X	X	X	X	X	X	X
7207-012 1/4-in. External Cartridge Tape Drive Model 012		X	X	X	X	X	X	X
7207-122 4 GB 1/4-in. External SLR5 QIC Tape Drive	X	X	X	X	X	X	X	X
7207-315 13 GB 1/4-in. External Cartridge Tape Drive		X	X	X	X	X	X	X
7208-011 External 8 mm Tape Drive		X	X	X	X	X	X	X
7208-341 20 Gb External 8 mm Tape		X	X	X	X	X	X	X
7208-345 60 GB External 8 mm Tape	X	X	X	X	X	X	X	X
7331-205 External 8 mm Tape Library		X	X	X	X	X	X	X
7331-305 External 8 mm Tape Library		X	X	X	X	X	X	X
7332-005 4 mm DDS-2 Tape Autoloader		X	X	X	X	X	X	X
7332-110 4 mm DDS-3 Tape Autoloader	X	X	X	X	X	X	X	X
7332-220 4 mm DDS-3 Tape Autoloader	X	X	X	X	X	X	X	X
7334-410 8 mm Tape Library	X	X	X	X	X	X	X	X
7336-205 4 mm Tape Library		X	X	X	X	X	X	X
7337-305 DLT Library		X	X	X	X	X	X	X
7337-306 DLT Library	X	X	X	X	X	X	X	X
7337-360 DLT Library	X	X	X	X	X	X	X	X
9348-012 Magnetic Tape Unit Model 012	X					X	X	X
Expansion Cabinets								
2032-C36 McDATA Cabinet for Directors and Fabric Switch	X	X	X	X		X	X	X
2042-C40 Inrange 40U Cabinet	X	X	X	X		X	X	X
2109-C36 TotalStorage San Cabinet	X	X						
7012-G02 Expansion Cabinet			X	X				
7014-S00 1.6 Meter System Rack - 32U	X	X	X	X				
7014-T00 2.0 Meter System Rack - 36U	X	X	X	X	X	X	X	X

External supported devices	Machine models							
	6 M 2	6 5 1	6 7 1	6 8 1	1 5 0	1 7 0	2 6 5	2 7 5
7014-T42 1.6 Meter System Rack - 42U	X	X	X	X	X	X	X	X
7015-R00 1.6 Meter System Rack - 32U	X	X	X	X				
7040-61D I/O Drawer			X	X				
7040-61R System Rack			X	X				
7311-D10 I/O Drawer	X							



Customer installation matrix and processor groups

This appendix includes customer installation matrix and processor group information for the following pSeries models:

- ▶ 7026 Model B80
- ▶ 7028 Model All
- ▶ 7029 Model All
- ▶ 7039 Model 651
- ▶ 7040 Model 671
- ▶ 7040 Model 681
- ▶ 7043 Model 150
- ▶ 7044 Model 170
- ▶ 7311 Model All
- ▶ 9112 Model 265
- ▶ 7014 Racks

Table F-1 on page 948 shows which models are intended to initially be set up by the customer and those that 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 F-1 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
7026	D5	B80	Yes	2856, 4362, 4363	All other features
7028	D5	All	Yes	None	All features
7029	D5	All	Yes	None	All features
7038	F5	6M2	No	All features	None
7039	G5	651	No	All features	None
7040	G5	671	No	All features	None
7040	H5	681	No	All features	None
7043	D5	150	Yes	6309	All other features
7044	D5	170	Yes	6309	All other features
7311		All	No	All features	None
9112	D5	265	Yes	None	All features
7014			Check with your IBM representative		



pSeries performance

In this section, we will discuss general information on the performance benchmarks.

Performance of pSeries systems

The good performance of a system is a relative concept because every individual has a different perception of it. System performance usually can be defined by:

- ▶ The response time for interactive users
- ▶ The complete time for batch jobs
- ▶ The number of reports finished each day
- ▶ The time a system needs for recovery after a failure
- ▶ No complaints about poor performance
- ▶ The graphical image can be redrawn a certain amount of times every second
- ▶ The system never needs changing to fix poor performance

More generally, a system displays good performance if it meets its performance requirements. Therefore, a clear definition of the performance requirements is really important before optimizing system performance, both in business terms of what the end users of the system actually get and in technical terms with empirical data. Good performance is achieved in three phases:

1. The system is adequately sized.
2. The system is initially set up to yield the maximum performance from the resources available.
3. Regular monitoring and tuning is performed, including upgrading if necessary.

Performance benchmarks

System performance is dependent on the kind of application software that is running on the system. Moreover, benchmarks are necessarily abstract and simplified models of all those environments. For this reason, benchmarks represent a good measure for comparing different systems rather than a precise tool for capacity planning for a given customer application environment.

No benchmark can fully characterize the performance of a system in a real production environment, because the behavior of benchmark applications is essentially constant on a given system and benchmarks are executed under ideal circumstances. Also, benchmarks are often specified algorithmically, or do not place restrictions on the amount of tuning that can be performed. In such situations, the application programming skill that is available to a vendor can play an extremely important role in determining the performance measured for the benchmark.

For a quick reference on performance, see the *IBM @server pSeries and RS6000 Performance Report*, found at:

http://www.ibm.com/servers/eserver/pseries/hardware/system_perf.pdf

Figure G-1 shows some examples of pSeries systems performance in the report provided at the previous URL.

Section 1 - SPEC2000 and Linpack Performance											
Model	Processor/ #CPUs	MHz	L1 Cache (KB)	L2 Cache (MB)	SPEC int_ 2000	SPEC base 2000	SPEC fp_ 2000	SPEC base 2000	DP	Linpack TPP	HPC
#43P-140u	604e	233	32/32	1.0	--	--	--	--	22.6	156.2	--
#43P-140n	604e	233	32/32	1.0	--	--	--	--	56	156.2	--
#43P-140	604e	332	32/32	1.0	--	--	--	--	59.9	179.7	--
43P-150	604e	250	32/32	1.0	105	99.4	90.8	90.8	43	170	--
43P-150	604e	375	32/32	1.0	--	--	--	--	64.8	255.7	--
44P-170	P3-II	333	32/64	1.0	202	196	277	274	363	833	--
44P-170	P3-II	400	32/64	4.0	280	271	359	355	461	1,052	--
44P-170	P3-II	450	32/64	8.0	346	333	434	426	503	1,440	--
#43P-260	P3/1	200	32/64	4.0	--	--	--	180	--	--	--
44P-270	P3-II/1	375	32/64	4.0	262	239	366	313	426	1,109	--
44P-270	P3-II/2	375	32/64	4.0	--	--	--	--	--	--	2,270
44P-270	P3-II/4	375	32/64	4.0	--	--	--	--	--	--	4,530
44P-270	P3-II/1	375	32/64	8.0	273	247	378	327	426	1,234	--

Figure G-1 Example of performance report

It provides the SPEC, TPC, Notesbench, and Relative OLTP, and other results of current IBM UNIX servers in a single summary.

System Performance Evaluation Corporation (SPEC)

The Standard Performance Evaluation Corporation (SPEC) is a non-profit corporation formed to establish, maintain, and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. The founders of this organization believe that the user community will benefit greatly from an objective series of applications-oriented tests, which can serve as common reference points and be considered during the evaluation process. While no benchmark can fully characterize overall system performance, the results of variety of realistic benchmarks can give valuable insight into expected real performance.

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, SPECint2000) are used for comparing the ability of a computer to complete single tasks. The SPEC rate

metrics (for example, SPECint_rate2000) measure the throughput or rate of a machine carrying out a number of tasks.

SPEC benchmarks include SPEC2000, SPEC SFS97, SPEC JVM98, and SPEC web99.

For more information of SPEC, visit the following Web site:

<http://www.spec.org>

Transaction oriented benchmarks

The Transaction Processing Council (TPC) was founded to define transaction processing and database benchmarks. It also was charged with delivering objective and verifiable performance data to the industry. TPC is a non-profit corporation of presently more than 40 hardware and software vendors, user organizations, and market research companies. As TPC benchmarks are focused on the overall performance of a system in a transaction-oriented environment, the TPC numbers will be used for comparing computers in a commercial environment.

The actual benchmarks are TPC-C, TPC-H, TPC-R, and TPC-W. Be advised that you do not need to compare different TPC benchmarks with each other, as the benchmarks are completely different and do not compare different major versions of the same benchmark.

For additional information, such as the versions of the published benchmarks, check the Transaction Processing Council's Web page at:

<http://www.tpc.org>

Relative Performance (rPerf)

The Relative Performance (rPerf) is an estimate of commercial processing performance. It is derived from an IBM analytical model that uses characteristics from IBM internal workloads, TPC, and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations, such as CPU, cache, and maximum memory available. However, the model does not simulate disk or network I/O operations. Although the model uses general database, and operating system parameters, it does not reflect specific database or AIX version or releases.

Unless otherwise indicated, rPerf is estimated only at the time a system is introduced. The pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to compare estimated IBM UNIX commercial


processing performance, actual system performance may vary and is dependent upon many factors, including system hardware configuration and software design and configuration. IBM withdrew Relative OLTP (ROLTP), and this figure should never be used to compare servers.

NotesBench benchmark

The Lotus NotesBench is a collection of benchmarks (workloads) for evaluating the performance of Domino servers. It measures many workloads by emulating the traffic that LAN-attached clients would generate when executing these workloads. The workloads (also called tests in the Lotus NotesBench for Domino R6 user guide) are software components that simulate the behavior of Domino workstation-to-server or server-to-server operations. They return measurements that let you evaluate server performance. The NotesBench is only available to hardware vendors and Lotus Business Partners who have fulfilled the prerequisite NotesBench training. NotesBench is not available to customers.

For additional information, check out the NotesBench Consortium Web site at:

<http://www.notesbench.org>



System life time information

This appendix contains the product life cycle dates of the pSeries and RS/6000 servers.

7006 IBM RS/6000 server

Table H-1 provides the system life time information for the 7006 RS/6000 server.

Table H-1 7006 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7006-41T	1994/05/24	1994/06/03	1997/01/10	-	7006-41W
7006-41W	1994/05/24	1994/06/03	1997/07/18	-	7043-140
7006-42T	1995/06/19	1995/07/07	1997/09/24	-	7043-140
7006-42W	1995/06/19	1995/07/07	1997/09/24	-	7043-240

7007 IBM RS/6000 server

Table H-2 provides the system life time information for the 7007-N40 RS/6000 server.

Table H-2 7007-N40 IBM RS/6000 Model N40 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7007-N40	1994/03/08	1994/03/25	1995/09/19	1997/06/30	-

7008 IBM RS/6000 server

Table H-3 provides the system life time information for the 7008 RS/6000 server.

Table H-3 7008 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7008-M20	1993/02/02	1993/03/26	1995/01/06	2001/12/31 -	-
7008-M2A	1993/02/02	1993/03/26	1994/10/26	2001/12/31	-

7009 IBM RS/6000 compact server

Table H-4 provides the system life time information for the 7009 RS/6000 server.

Table H-4 7009 IBM RS/6000 compact server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7009-C10	1994/05/24	1994/06/03 1	1997/07/18	-	7043-140
7009-C20	1995/06/19	1995/07/07	1998/01/30	-	None

7010 IBM RS/6000 server

Table H-5 provides the system life time information for the 7010 RS/6000 server.

Table H-5 7010 IBM RS/6000 server

Type model	Announced	Available marketing	Withdrawn	Service discontinued	Replaced by
7010-160	1995/02/07	1995/02/24	1996/05/20	2004/12/31	-
7010-150	1992/02/02	1993/03/26	1996/05/20	2004/12/31	-
7010-140	1993/09/21	1993/10/15	1996/05/20	2004/12/31	-
7010-130	1991/03/12	1991/04/26	1995/01/06	2004/12/31	7010-140
7010-120	1990/02/15	1990/04/15	1993/08/18	2004/12/31	7010-130

7011 IBM RS/6000 server

Table H-6 provides the system life time information for the 7011 RS/6000 server.

Table H-6 7011 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7011-220	1992/01/21	1992/04/24	1995/01/06	-	7011-250
7011-22W	1992/01/21	1992/04/24	1995/01/06	-	7011-25W
7011-230	1993/05/18	1993/05/28	1995/01/06	-	7011-250
7011-23S	1993/05/18	1993/05/28	1994/10/26	-	7011-25S

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7011-23T	1993/05/18	1993/05/28	1994/10/26	-	7011-25T
7011-23W	1993/05/18	1993/05/28	1994/10/26	-	7011-25W
7011-250	1993/09/21	1993/10/15	1997/07/18	-	7043-140
7011-25S	1993/09/21	1993/10/15	1996/10/25	-	7011-250
7011-25T	1993/09/21	1993/10/15	1997/07/18	-	7043-140
7011-25W	1993/09/21	1993/10/15	1996/10/25	-	7011-25T

7012 IBM RS/6000 server

Table H-7 provides the system life time information for the 7012 RS/6000 server.

Table H-7 7012 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7012-320	1990/02/15	1990/06/30	1992/10/28	-	7012-32H
7012-32H	1991/03/12	1991/03/29	1994/10/26	-	7009-C10
7012-340	1992/01/21	1992/02/28	1994/11/04	-	7012-370
7012-34H	1993/07/13	1993/08/06	1994/10/26	-	None
7012-350	1992/01/21	1992/02/28	1993/08/18	-	7012-360
7012-355	1993/02/02	1993/02/26	1994/10/26	-	7006-41T
7012-360	1993/02/02	1993/02/26	1994/11/04	-	7012-370
7012-365	1993/02/02	1993/02/26	1994/10/26	-	7006-41T
7012-36T	1993/05/18	1993/06/04	1994/10/26	2001/12/31	7006-41T
7012-370	1993/02/02	1993/02/26	1996/05/20	-	7012-390
7012-375	1993/02/02	1993/02/26	1994/10/26	-	7006-41T
7012-37T	1993/05/18	1993/06/04	1996/05/20	-	7030-3BT
7012-380	1994/05/24	1994/06/10	1996/05/20	-	7012-390
7012-390	1994/05/24	1994/06/10	1997/07/18	-	7012-39H

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7012-39H	1995/02/07	1995/02/17	1998/01/30	-	7012-397
7012-397	1997/10/06	1997/10/31	1999/03/19	-	7043-260
7012-G30	1994/10/04	1994/12/23	1996/10/23	-	7012-G40
7012-G40	1996/07/23	1996/08/30	1998/01/08	-	7025-F50

7013 IBM RS/6000 server

Table H-8 provides the system life time information for the 7013 RS/6000 server.

Table H-8 7013 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7013-520	1990/02/15	1990/06/00	1992/04/21	-	7013-52H
7013-52H	1992/01/21	1995/02/07	1995/01/06	-	7013-570
7013-530	1990/02/15	1990/06/00	1992/01/02	-	7013-53H
7013-53H	1991/10/02	1991/10/04	1993/08/18	-	7013-570
7013-540	1990/02/15	1990/09/00	1992/01/02	-	7013-550
7013-550	1990/10/30	1991/03/29	1993/08/18	-	7013-580
7013-55L	1993/05/18	1993/06/04	1994/10/26	-	7013-570
7013-560	1992/01/21	1992/03/27	1993/12/21	-	7013-580
7013-570	1993/02/02	1993/02/19	1996/05/20	-	7013-590
7013-580	1992/09/22	1993/05/21	1996/05/20	-	7013-J30
7013-58H	1993/09/21	1993/10/22	1996/10/25	-	7013-590
7013-590	1993/09/21	1993/10/22	1997/09/24	-	7013-595
7013-59H	1994/05/24	1994/06/10	1997/01/10	-	7013-595
7013-591	1995/07/25	1995/08/18	1997/07/18	-	7013-595
7013-595	1996/10/08	1996/11/08	1999/01/08	-	-
7013-J30	1994/10/04	1994/12/23	1996/10/23	-	7013-J40

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7013-J40	1996/07/23	1996/08/30	1998/01/08	-	7013-J50
7013-J50	1997/04/15	1997/04/30	1999/01/08	-	-
7013-S70	1997/10/06	1997/12/05	1999/12/13	-	-
7013-S7A	1998/10/05	1998/10/23	2000/12/01	-	7017-S80

7014 IBM RS/6000 server

Table H-9 provides the system life time information for the 7014 RS/6000 server.

Table H-9 7014 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7014-S00	1997/10/06	1997/10/31	2001/07/17	-	None
7014-T00	2000/05/09	2000/06/09	-	-	-
7014-T42	2000/05/09	2000/08/04	-	-	-

7015 IBM RS/6000 server

Table H-10 provides the system life time information for the 7015 RS/6000 server.

Table H-10 7015 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7015-930	1990/02/15	1990/07/00	1992/07/15	-	-
7015-950	1991/05/07	1991/06/28	1993/12/21	-	-
7015-970	1992/04/21	1992/06/26	1993/08/18	-	-
7015-97B	1993/02/02	1993/03/05	1995/01/06	-	7015-990
7015-980	1992/09/22	1992/10/23	1993/08/18	-	-
7015-98B	1993/02/02	1993/03/05	1996/05/20	-	7015-R24
7015-990	1993/09/21	1993/10/29	1996/05/20	-	7015-99K

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7015-R00	1994/05/24	1994/06/10	2001/07/17	-	-
7015-R10	1994/05/24	1994/06/10	1996/05/20	-	7015-R20
7015-R20	1994/05/24	1994/06/10	1998/01/30	-	7015-R50
7015-R21	1995/07/25	1995/08/18	1996/10/25	-	7015-R20
7015-R24	1994/05/24	1994/07/15	1998/01/30	-	7015-R50
7015-R30	1994/10/04	1994/12/23	1996/10/23	-	7015-R40
7015-R40	1996/07/23	1996/08/30	1998/01/08	-	7015-R50
7015-R50	1997/04/15	1997/04/30	2000/08/15	-	7026-M80
7015-S70	1997/10/06	1997/10/31	1999/12/13	-	-
7015-S7A	1998/10/05	1998/10/23	2000/12/01	-	7017-S80

7017 IBM RS/6000 enterprise server

Table H-11 provides the system life time information for the 7017 RS/6000 enterprise server.

Table H-11 7017 IBM RS/6000 enterprise server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7017-S70	1997/10/06	1997/10/31	1999/12/13	-	7017-S7A
7017-S7A	1998/10/05	1998/10/23	2000/12/01	-	7017-S80
7017-S80	1999/09/13	1999/09/24	2001/08/31	-	7017-S85
7017-S85	2000/10/03	2000/11/17	2003/03/28	-	7040-681

7024 IBM RS/6000 server

Table H-12 provides the system life time information for the 7024 RS/6000 server.

Table H-12 7024 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7024-E20	1995/10/10	1995/10/20	1997/07/18	-	7024-E30
7024-E30	1996/04/23	1996/05/03	1999/03/19	-	7025-F40

7025 IBM RS/6000 server

Table H-13 provides the system life time information for the 7025 RS/6000 server.

Table H-13 7025 IBM RS/6000 server

Type model	Announced	Available marketing	Withdrawn	Service discontinued	Replaced by
7025-F30	1996/02/20	1996/03/01	1998/01/08	-	7025-F40
7025-F40	1996/10/08	1996/11/08	2000/05/08	-	None
7025-F50	1997/04/15	1997/04/25	2001/07/17	-	None
7025-H70	1999/09/13	1999/09/24	2001/07/17 -		None
7025-F80	2000/05/09	2000/06/09	2001/07/13	-	7025-6F1
7025-6F1	2001/04/17	2001/04/27	2003/09/12	-	7028-6E4
7025-6F0	2001/06/05	2001/06/15	2003/09/12	-	7028-6E4

7026 IBM RS/6000 server

Table H-14 provides the system life time information for the 7026 RS/6000 server.

Table H-14 7026 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7026-H10	1996/10/08	1997/02/14	1998/02/27	-	None
7026-H50	1998/02/09	1998/02/20	2000/12/01	-	7026-H70
7026-H70	1999/04/06	1999/04/23	2001/07/17	-	None
7026-H80	2000/05/09	2000/06/09	2001/07/13	-	7026-6H1
7026-M80	2000/05/09	2000/06/09	2002/01/31	-	7026-6H1
7026-B80	2000/10/03	2000/10/27	2003/12/12	-	7028-6C1
7026-6H1	2001/04/17	2001/04/27	2003/09/12	-	7028-6C4
7026-6H0	2001/06/05	2001/06/15	2003/09/12	-	7028-6C4
7026-6M1	2001/09/04	2001/09/21	2003/09/12	-	7038-6M2

7027 IBM high capacity drawer

Table H-15 provides the system life time information for the 7027 IBM high capacity drawer.

Table H-15 7027 IBM high capacity drawer

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7027-HSC	1996/04/23	1996/05/24	1998/08/21 -	-	7133-XXX
7027-HSD	1996/04/23	1996/05/24	1998/08/21	-	7133-XXX

7028 IBM @server pSeries 610 server

Table H-16 provides the system life time information for the 7026 IBM @server pSeries 610 server.

Table H-16 7028 IBM @server pSeries 610 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7028-6C1	2001/10/04	2001/10/26	-	-	-
7028-6E1	2001/10/04	2001/10/26	-	-	-
7028-6C4	2002/06/25	2002/08/30	-	-	-
7028-6E4	2002/06/25	2002/08/30	-	-	-

7030 IBM RS/6000 server

Table H-17 provides the system life time information for the 7030 RS/6000 server.

Table H-17 7030 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7030-3AT	1994/05/24	1994/06/10	1997/01/10	2004/06/30	7012-390
7030-3BT	1994/05/24	1994/08/12	1998/01/08	2004/06/30	7012-397
7030-3CT	1995/02/07	1995/02/17	1998/01/08	2004/06/30	7012-397
7030-397	1997/10/06	1997/10/31	1999/01/19	-	7043-260

7038-6M2 IBM @server pSeries 650

Table H-18 provides the system life time information for the 7038 IBM @server pSeries 650.

Table H-18 7038-6M2 IBM @server pSeries 650 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7038-6M2	2002/11/12	2002/12/06	-	-	-

7039-651 IBM @server pSeries 655 Server Model 651

Table H-19 provides the system life time information for the 7038 IBM @server pSeries 655 Model 651.

Table H-19 7039-651 IBM @server pSeries 655 Model 651

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7039-651	2002/11/12	2002/12/13	-	-	-

7040 IBM @server pSeries 670/690 servers

Table H-20 provides the system life time information for the 7040 IBM @server pSeries 670/690 servers.

Table H-20 7040 IBM @server pSeries 670/690 servers

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7040-671	2002/04/09	2002/04/26	-	-	-
7040-681	2001/10/04	2001/12/14	-	-	-
7040-61D	2001/10/04	2001/12/14	-	-	-
7040-61R	2001/10/04	2001/12/14	-	-	-
7040-W42	2002/11/12	2002/12/13	-	-	-

7043 IBM RS/6000 43P server

Table H-21 provides the system life time information for the 7043 RS/6000 server.

Table H-21 7043 IBM RS/6000 43P server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7043-140	1996/10/08	1996/11/08	2000/12/01	-	7043-150
7043-150	1998/10/05	1998/10/23	2003/12/12	-	7028-6E1
7043-240	1996/10/08	1996/11/08	1999/03/19	-	7043-260

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7043-260	1998/10/05	1998/10/23	2000/12/01	-	7044-270
7043-270	2000/02/07	2000/02/25	2003/09/12	-	None

7044 IBM RS/6000 server

Table H-22 provides the system life time information for the 7044 RS/6000 server.

Table H-22 7044 IBM RS/6000 server

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7044-170	2000/02/07	2000/02/25	2003/12/12	-	7028-6E1
7044-270	2000/02/07	2000/02/25	2003/09/12	-	7029-6E3

7046-B50 IBM RS/6000 server

Table H-23 provides the system life time information for the 7044 RS/6000 Model B50.

Table H-23 7046-B50 IBM RS/6000 Model B50

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7046-B50	1999/09/13	1999/09/24	2003/09/12	-	None

7202-900 IBM RS/6000 expansion rack Model 900

Table H-24 provides the system life time information for the 7202-900 RS/6000 Expansion Rack Model 900.

Table H-24 7202-900 IBM RS/6000 expansion rack Model 900

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7202-900	1990/01/23	1990/03/30	1996/05/20	-	7015-R00

7236-001 IBM RS/6000 Media Streamer Model 001

Table H-25 provides the system life time information for the 7236-001 IBM RS/6000 Media Streamer Model 001.

Table H-25 7236-001 IBM RS/6000 Media Streamer Model 001

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7236-001	1996/12/03	1997/02/28	1998/06/30	2001/12/31	None

7248 IBM RS/6000 43P series

Table H-26 provides the system life time information for the 7248 IBM RS/6000 43P series.

Table H-26 7248 IBM RS/6000 43P series

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7248-100	1995/06/19	1995/07/07	1997/01/10	2003/12/31	7248-120
7248-120	1995/06/19	1995/07/07	1997/07/18	2003/12/31	7043-140
7248-132	1995/06/19	1995/07/07	1997/07/18	2003/12/31	7043-140

7311 IBM Rack-mounted Drawers

Table H-27 provides the system life time information for the 7311 IBM Rack-mounted Drawers.

Table H-27 7311 IBM Rack-Mounted Drawers

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7311-D10	2002/11/12	2002/12/06	-	-	-
7311-D20	2002/11/12	2002/12/06	-	-	-

7315-C01 IBM Hardware Management Console

Table H-28 provides the system life time information for the 7315-C01 IBM Hardware Management Console C01.

Table H-28 7315-C01 IBM Hardware Management Console C01

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7315-C01	2002/10/08	2002/10/25	2003/11/14	-	7315-C02
7315-C02	2003/06/24	2003/07/18	2003/11/18	-	7315-C03
7315-C03	2003/11/18	2003/11/28	-	-	-
7315-CR2	2003/09/09	2003/09/19	-	-	-

7316 IBM Rack-mounted Console

Table H-29 provides the system life time information for the 7316 IBM Rack-Mounted Console.

Table H-29 7316 IBM Rack-Mounted Console

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7316-TF1	2000/12/05	2001/02/09	2002/09/13	-	316-TF2
7316-TF2	2002/08/13	2002/09/13	-	-	-

7317 IBM RS/6000 Telecom Server and External SCSI Disk

Table H-30 provides the system life time information for the 7317 IBM RS/6000 Telecom Server and External SCSI Disk.

Table H-30 7317 IBM RS/6000 Telecom Server and External SCSI Disk

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7317-D10	1996/10/08	1996/11/15	1999/12/13	-	none
7317-F3L	1996/10/08	1996/11/15	1999/12/13	-	none

7318 IBM Serial Communications Network Server

Table H-31 provides the system life time information for 7318 IBM Serial Communications Network Server.

Table H-31 7318 IBM Serial Communications Network Server


Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7318-P10	1994/05/24	1994/06/10	1998/12/31	-	-
7318-S20	1994/05/24	1994/06/10	1998/12/31	-	-

7319 IBM RS/6000 Fibre Channel Switch and Adapter

Table H-32 provides the system life time information for the 7319 IBM RS/6000 Fibre Channel Switch and Adapter.

Table H-32 7319 IBM RS/6000 Fibre Channel Switch and Adapter

Type model	Announced	Available	Marketing withdrawn	Service discontinued	Replaced by
7319-100 1	994/05/24	1994/07/29	1996/05/20	1999/12/31	7319-110
7319-110 1	995/10/10	1996/03/29	1997/09/24	1999/12/31	None



Firmware

This appendix contains the link to download the newest firmware for system, adapter and devices, which is:

<http://techsupport.services.ibm.com/server/mdownload/download.html>

Once you select to download a piece of microcode, you will be presented with a license agreement that you must accept before you will be allowed to complete the download. You could be asked for a password on some files; the eleven character case sensitive password is ****RS/6000****.

To further understand each microcode package, including how to determine the level of microcode already installed, see the description for each microcode package. For instructions on how to download these microcode updates, refer to the microcode download procedures. What's New provides information on what has recently changed on this page.

Abbreviations and acronyms

ABI	Application Binary Interface	AUI	Attached Unit Interface
ABM	Advanced Battery Management	AWT	Abstract Window Toolkit
AC	Alternating Current	BCT	Branch on CounT
ACF	Automatic Cartridge Feed	BFF	Backup File Format
ACI	Access Control Information	BI	Business Intelligence
ACL	Access Control List	BIND	Berkeley Internet Name Domain
ADSM	ADSTAR Distributed Storage Manager	BIST	Built-In Self-Test
ADSTAR	Advanced Storage and Retrieval	BLAS	Basic Linear Algebra Subprograms
AFPA	Adaptive Fast Path Architecture	BLOB	Binary Large Object
AFS®	Andrew File System	BLV	Boot Logical Volume
AH	Authentication Header	BOOTP	Boot Protocol
AIO	Asynchronous I/O	BOS	Base Operating System
AIX	Advanced Interactive Executive	BPA	Bulk Power Assembly
ANSI	American National Standards Institute	BPC	Bulk Power Controller
APAR	Authorized Program Analysis Report	BPC	Bulk Power Controller
API	Application Programming Interface	BPF	Berkeley Packet Filter
AppA	Application Audio	BPR	Bulk Power Regulator
AppV	Application Video	BSC	Binary Synchronous Communications
ARP	Address Resolution Protocol	BSD	Berkeley Software Distribution
ASCI	Accelerated Strategic Computing Initiative	CA	Certificate Authority
ASCII	American National Standards Code for Information Interchange	CAD	Computer-Aided Design
ASR	Address Space Register	CAE	Computer-Aided Engineering
ATM	Asynchronous Transfer Mode	CAM	Computer-Aided Manufacturing
AuditRM	Audit Log Resource Manager	CATE	Certified Advanced Technical Expert
		CATIA	Computer-Graphics Aided Three-Dimensional Interactive Application
		CBC	Cipher Block Chain

CCA	Common Cryptographic Architecture	CRL	Certificate Revocation List
CCM	Common Character Mode	CRM	Concurrent Resource Manager
CD	Compact Disk	CSID	Character Set ID
CDE	Common Desktop Environment	CSM	Cluster Systems Management
CDLI	Common Data Link Interface	CSR	Customer Service Representative
CD-R	CD Recordable	CSS	Communication Subsystems Support
CD-ROM	Compact Disk-Read Only Memory	CSU	Customer Set-Up
CE	Customer Engineer	CTQ	Command Tag Queuing
CEC	Central Electronics Complex	CUoD	Capacity Upgrade on Demand
CFD	Computational Fluid Dynamics	CWOF	Cascading without Fallback
CFM	Configuration File Manager	CWR	Congestion Window Reduced
CGE	Common Graphics Environment	CWS	Control Workstation
CHRP	Common Hardware Reference Platform	DAA	Dual Active Accessor
CIFS	Common Internet File System	DAD	Duplicate Address Detection
CIM	Common Information Model	DAS	Dual Attach Station
CISPR	International Special Committee on Radio Interference	DASD	Direct Access Storage Device
CIU	Core Interface Unit	DAT	Digital Audio Tape
CIU	Core Interface Unit	DBCS	Double Byte Character Set
CLI	Command Line Interface	DBE	Double Buffer Extension
CLIO/S	Client Input/Output Sockets	DC	Direct Current
CLVM	Concurrent LVM	DCA	Distributed Converter Assembly
CMOS	Complimentary Metal-Oxide Semiconductor	DCE	Distributed Computing Environment
CMP	Certificate Management Protocol	DCEM	Distributed Command Execution Manager
COFF	Common Object File Format	DCM	Dual Chip Module
COLD	Computer Output to Laser Disk	DCUoD	Dynamic Capacity Upgrade on Demand
CPU	Central Processing Unit	DDC	Display Data Channel
CRC	Cyclic Redundancy Check	DDR	Double Data Rate
		DDS	Digital Data Storage
		DE	Dual-Ended

DES	Data Encryption Standard	DSP	Digital Sound Processor
DFL	Divide Float	DSU	Data Service Unit
DFP	Dynamic Feedback Protocol	DTD	Document Type Definition
DFS	Distributed File System	DTE	Data Terminating Equipment
DGD	Dead Gateway Detection	DVD	Digital Versatile Disk
DH	Diffie-Hellman	DVI	Digital Video Interface
DHCP	Dynamic Host Configuration Protocol	DW	Data Warehouse
DIMM	Dual In-Line Memory Module	DWA	Direct Window Access
DIN	Deutsche Industrie Norm connector	EA	Effective Address
DIP	Direct Insertion Probe	EC	Engineering Change
DIT	Directory Information Tree	ECC	Error Checking and Correcting
DIVA	Digital Inquiry Voice Answer	ECN	Explicit Congestion Notification
DLPAR	Dynamic LPAR	EEH	Extended Error Handling
DLT	Digital Linear Tape	EEPROM	Electrically Erasable Programmable Read Only Memory
DMA	Direct Memory Access	EFI	Extensible Firmware Interface
DMT	Directory Management Tool	EHD	Extended Hardware Drivers
DMTF	Distributed Management Task Force	EIA	Electronic Industries Association
DN	Distinguished Name	EIM	Enterprise Identity Mapping
DNLC	Dynamic Name Lookup Cache	EISA	Extended Industry Standard Architecture
DNS	Domain Naming System	ELA	Error Log Analysis
DOE	Department of Energy	ELF	Executable and Linking Format
DOI	Domain of Interpretation	EMEA	East Middle East Asia
DOM	Document Object Model	EMF	Electro-magnetic Frequency
DOS	Disk Operating System	EMIF	Multiple Image Facility
DPCL	Dynamic Probe Class Library	EMU	European Monetary Union
DRAM	Dynamic Random Access Memory	EOF	End of File
DRM	Dynamic Reconfiguration Manager	EPOW	Environmental and Power Warning
DS	Differentiated Service	EPROM	Erasable Programmable Read-only Memory
DSA	Dynamic Segment Allocation		
DSE	Diagnostic System Exerciser		
DSMIT	Distributed SMIT		

ERRM	Event Response Resource Manager	FTP	File Transfer Protocol
ESCON	Enterprise System Connection	GAI	Graphic Adapter Interface
ESID	Effective Segment ID	GAMESS	General Atomic and Molecular Electronic Structure System
ESP	Encapsulating Security Payload	GBps	Gigabytes per second
ESSL	Engineering and Scientific Subroutine Library	Gbps	Gigabits per second
ETML	Extract, Transformation, Movement, and Loading	GID	Group ID
F/W	Fast and Wide	GPFS	General Parallel File System
FC	Feature Code	GPR	General-Purpose Register
FC	Fibre Channel	GUI	Graphical User Interface
FCAL	Fibre Channel Arbitrated Loop	GUID	Globally Unique Identifier
FCC	Federal Communication Commission	HACMP	High Availability Cluster Multi Processing
FCP	Fibre Channel Protocol	HACWS	High Availability Control Workstation
FDDI	Fiber Distributed Data Interface	HBA	Host Bus Adapters
FDPR	Feedback Directed Program Restructuring	HCON	IBM AIX Host Connection Program/6000
FDX	Full Duplex	HDX	Half Duplex
FIFO	First In/First Out	HFT	High Function Terminal
FIPS	Federal Information Processing Standards	HIPPI	High Performance Parallel Interface
FLASH EPROM	Flash Erasable Programmable Read-Only Memory	HiPS	High Performance Switch
FLIH	First Level Interrupt Handler	HiPS LC-8	Low-Cost Eight-Port High Performance Switch
FMA	Floating point Multiply Add operation	HMC	Hardware Management Console
FPR	Floating Point Register	HMT	Hardware Multithreading
FPU	Floating Point Unit	HostRM	Host Resource Manager
FRCA	Fast Response Cache Architecture	HP	Hewlett-Packard
FRU	Field Replaceable Unit	HPF	High Performance FORTRAN
FSRM	File System Resource Manager	HPSSDL	High Performance Supercomputer Systems Development Laboratory
		HP-UX	Hewlett-Packard UNIX
		HTML	Hyper-text Markup Language
		HTTP	Hypertext Transfer Protocol

Hz	Hertz	IPL	Initial Program Load
I/O	Input/Output	IPSec	IP Security
I²C	Inter Integrated-Circuit Communications	IrDA	Infrared Data Association (which sets standards for infrared support including protocols for data interchange)
IAR	Instruction Address Register		
IBF	Internal Battery Feature		
IBM	International Business Machines	IRQ	Interrupt Request
		IS	Integrated Service
ICCCM	Inter-Client Communications Conventions Manual	ISA	Industry Standard Architecture or Instruction Set Architecture
ICE	Inter-Client Exchange		
ICElib	Inter-Client Exchange library	ISAKMP	Internet Security Association Management Protocol
ICMP	Internet Control Message Protocol	ISB	Intermediate Switch Board
ID	Identification	ISDN	Integrated-Services Digital Network
IDE	Integrated Device Electronics		
IDL	Interface Definition Language	ISMP	InstallShield Multi-Platform
IDS	Intelligent Decision Server	ISNO	Interface Specific Network Options
IEEE	Institute of Electrical and Electronics Engineers	ISO	International Organization for Standardization
IETF	Internet Engineering Task Force	ISV	Independent Software Vendor
IHS	IBM HTTP Server	ITSO	International Technical Support Organization
IHV	Independent Hardware Vendor	IXFR	Incremental Zone Transfer
IIOIP	Internet Inter-ORB Protocol	JBOD	Just a Bunch of Disks
IJG	Independent JPEG Group	JCE	Java Cryptography Extension
IKE	Internet Key Exchange	JDBC	Java Database Connectivity
ILMI	Integrated Local Management Interface	JFC	Java Foundation Classes
		JFS	Journaled File System
ILS	International Language Support	JSSE	Java Secure Sockets Extension
IM	Input Method	JTAG	Joint Test Action Group
INRIA	Institut National de Recherche en Informatique et en Automatique	JVMPI	Java Machine Profiling Interface
		KDC	Key Distribution Center
IP	Internetwork Protocol (OSI)	L1	Level 1
IPAT	IP address takeover	L2	Level 2

L3	Level 3	Mbps	Megabits Per Second
LAM	Loadable Authentication Module	MBps	Megabytes Per Second
LAN	Local Area Network	MCA	Micro Channel® Architecture
LANE	Local Area Network Emulation	MCAD	Mechanical Computer-Aided Design
LAPI	Low-Level Application Programming Interface	MCM	Multichip Module
LDAP	Lightweight Directory Access Protocol	MDF	Managed Object Format
LDIF	LDAP Directory Interchange Format	MDI	Media Dependent Interface
LED	Light Emitting Diode	MES	Miscellaneous Equipment Specification
LFD	Load Float Double	MFLOPS	Million of Floating point Operations Per Second
LFT	Low Function Terminal	MIB	Management Information Base
LID	Load ID	MII	Media Independent Interface
LLNL	Lawrence Livermore National Laboratory	MIP	Mixed-Integer Programming
LMB	Logical Memory Block	MLR1	Multi-Channel Linear Recording 1
LP	Logical Partition	MMF	Multi-Mode Fibre
LP64	Long-Pointer 64	MODS	Memory Overlay Detection Subsystem
LPAR	Logical Partitioning	MP	Multiprocessor
LPI	Lines Per Inch	MPC-3	Multimedia PC-3
LPP	Licensed Program Product	MPI	Message Passing Interface
LPR/LPD	Line Printer/Line Printer Daemon	MPIO	Multipath I/O
LRU	Least Recently Used	MPOA	Multiprotocol over ATM
LTG	Logical Track Group	MPP	Massively Parallel Processing
LV	Logical Volume	MPS	Mathematical Programming System
LVCB	Logical Volume Control Block	MSS	Maximum Segment Size
LVD	Low Voltage Differential	MST	Machine State
LVM	Logical Volume Manager	MTU	Maximum Transmission Unit
MAP	Maintenance Analysis Procedure	MWCC	Mirror Write Consistency Check
MASS	Mathematical Acceleration Subsystem	MX	Mezzanine Bus
MAU	Multiple Access Unit	NBC	Network Buffer Cache
MBCS	Multi-Byte Character Support	NCP	Network Control Point
		ND	Neighbor Discovery

NDP	Neighbor Discovery Protocol	PAG	Process Authentication Group
NDS	Novell Directory Services	PAM	Pluggable Authentication Mechanism
NFB	No Frame Buffer	PAP	Privileged Access Password
NFS	Network File System	PBLAS	Parallel Basic Linear Algebra Subprograms
NHRP	Next Hop Resolution Protocol	PCB	Protocol Control Block
NIM	Network Installation Management	PCI	Peripheral Component Interconnect
NIS	Network Information Service	PDT	Paging Device Table
NL	National Language	PDU	Power Distribution Unit
NLS	National Language Support	PE	Parallel Environment
NT-1	Network Terminator-1	PEDB	Parallel Environment Debugging
NTF	No Trouble Found	PEX	PHIGS Extension to X
NTP	Network Time Protocol	PFS	Perfect Forward Security
NUMA	Non-Uniform Memory Access	PGID	Process Group ID
NUS	Numerical Aerodynamic Simulation	PHB	Processor Host Bridges
NVRAM	Non-Volatile Random Access Memory	PHY	Physical Layer
NWP	Numerical Weather Prediction	PID	Process ID
OACK	Option Acknowledgment	PIOFS	Parallel Input Output File System
OCS	Online Customer Support	PKCS	Public-Key Cryptography Standards
ODBC	Open DataBase Connectivity	PKI	Public Key Infrastructure
ODM	Object Data Manager	PKR	Protection Key Registers
OEM	Original Equipment Manufacturer	PMTU	Path MTU
OLAP	Online Analytical Processing	POE	Parallel Operating Environment
OLTP	Online Transaction Processing	POP	Power-On Password
ONC+	Open Network Computing	POSIX	Portable Operating Interface for Computing Environments
OOUI	Object-Oriented User Interface	POST	Power-On Self-test
OSF	Open Software Foundation, Inc.	POWER	Performance Optimization With Enhanced RISC (Architecture)
OSL	Optimization Subroutine Library	PPC	PowerPC
OSLp	Parallel Optimization Subroutine Library	PPM	Piecewise Parabolic Method
P2SC	POWER2 Single/Super Chip		

PPP	Point-to-Point Protocol	RISC	Reduced Instruction-Set Computer
PREP	PowerPC Reference Platform®	RMC	Resource Monitoring and Control
PRNG	Pseudo-Random Number Generator	ROLTP	Relative Online Transaction Processing
PSE	Portable Streams Environment	RPA	RS/6000 Platform Architecture
PSSP	Parallel System Support Program	RPC	Remote Procedure Call
PTF	Program Temporary Fix	RPL	Remote Program Loader
PTPE	Performance Toolbox Parallel Extensions	RPM	Red Hat Package Manager
PTX	Performance Toolbox	RSC	RISC Single Chip
PV	Physical Volume	RSCT	Reliable Scalable Cluster Technology
PVC	Permanent Virtual Circuit	RSE	Register Stack Engine
PVID	Physical Volume Identifier	RSVP	Resource Reservation Protocol
QMF™	Query Management Facility	RTC	Real-Time Clock
QoS	Quality of Service	RVSD	Recoverable Virtual Shared Disk
QP	Quadratic Programming	SA	Secure Association
RAID	Redundant Array of Independent Disks	SACK	Selective Acknowledgments
RAM	Random Access Memory	SAN	Storage Area Network
RAN	Remote Asynchronous Node	SAR	Solutions Assurance Review
RAS	Reliability, Availability, and Serviceability	SAS	Single Attach Station
RDB	Relational DataBase	SASL	Simple Authentication and Security Layer
RDBMS	Relational DataBase Management System	SBCS	Single-Byte Character Support
RDF	Resource Description Framework	ScaLAPACK	Scalable Linear Algebra Package
RDISC	ICMP Router Discovery	SCB	Segment Control Block
RDN	Relative Distinguished Name	SCSI	Small Computer System Interface
RDP	Router Discovery Protocol	SCSI-SE	SCSI-Single Ended
RFC	Request for Comments	SDK	Software Development Kit
RIO	Remote I/O	SDLC	Synchronous Data Link Control
RIP	Routing Information Protocol	SDR	System Data Repository
RIPL	Remote Initial Program Load		

SDRAM	Synchronous Dynamic Random Access Memory	SPI	Security Parameter Index
SE	Single Ended	SPM	System Performance Measurement
SEPBU	Scalable Electrical Power Base Unit	SPOT	Shared Product Object Tree
SGI	Silicon Graphics Incorporated	SPS	SP Switch
SGID	Set Group ID	SPS-8	Eight-Port SP Switch
SHLAP	Shared Library Assistant Process	SRC	System Resource Controller
SID	Segment ID	SRN	Service Request Number
SIT	Simple Internet Transition	SSA	Serial Storage Architecture
SKIP	Simple Key Management for IP	SSC	System Support Controller
SLB	Segment Lookaside Buffer	SSL	Secure Socket Layer
SLIH	Second Level Interrupt Handler	STFDU	Store Float Double with Update
SLIP	Serial Line Internet Protocol	STP	Shielded Twisted Pair
SLR1	Single-Channel Linear Recording 1	SUID	Set User ID
SM	Session Management	SUP	Software Update Protocol
SMB	Server Message Block	SVC	Switch Virtual Circuit
SMIT	System Management Interface Tool	SVC	Supervisor or System Call
SMP	Symmetric Multiprocessor	SWVPD	Software Vital Product Data
SMS	System Management Services	SYNC	Synchronization
SNG	Secured Network Gateway	TCB	Trusted Computing Base
SNIA	Storage Networking Industry Association	TCE	Translate Control Entry
SNMP	Simple Network Management Protocol	Tcl	Tool Command Language
SOI	Silicon-on-Insulator	TCP/IP	Transmission Control Protocol/Internet Protocol
SP	IBM RS/6000 Scalable POWER Parallel Systems	TCQ	Tagged Command Queuing
SP	Service Processor	TGT	Ticket Granting Ticket
SPCN	System Power Control Network	TLB	Translation Lookaside Buffer
SPEC	System Performance Evaluation Corporation	TLS	Transport Layer Security
		TOS	Type Of Service
		TPC	Transaction Processing Council
		TPP	Toward Peak Performance
		TSE	Text Search Engine
		TTL	Time To Live
		UCS	Universal Coded Character Set

UDB EEE	Universal Database and Enterprise Extended Edition	WBEM	Web-based Enterprise Management
UDF	Universal Disk Format	WLM	Workload Manager
UDI	Uniform Device Interface	WTE	Web Traffic Express
UIL	User Interface Language	XCOFF	Extended Common Object File Format
ULS	Universal Language Support	XIE	X Image Extension
UNI	Universal Network Interface	XIM	X Input Method
UP	Uniprocessor	XKB	X Keyboard Extension
USB	Universal Serial Bus	XL F	XL Fortran
USLA	User-Space Loader Assistant	XML	Extended Markup Language
UTF	UCS Transformation Format	XOM	X Output Method
UTM	Uniform Transfer Model	XPM	X Pixmap
UTP	Unshielded Twisted Pair	XSSO	Open Single Sign-on Service
UUCP	UNIX-to-UNIX Communication Protocol	XTF	Extended Distance Feature
VACM	View-based Access Control Model	XVFB	X Virtual Frame Buffer
VESA	Video Electronics Standards Association		
VFB	Virtual Frame Buffer		
VG	Volume Group		
VGDA	Volume Group Descriptor Area		
VGSA	Volume Group Status Area		
VHDCI	Very High Density Cable Interconnect		
VIPA	Virtual IP Address		
VLAN	Virtual Local Area Network		
VMM	Virtual Memory Manager		
VP	Virtual Processor		
VPD	Vital Product Data		
VPN	Virtual Private Network		
VSD	Virtual Shared Disk		
VSM	Visual System Manager		
VSS	Versatile Storage Server		
VT	Visualization Tool		
WAN	Wide Area Network		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 991. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765
- ▶ *AIX and Linux Interoperability*, SG24-6622
- ▶ *The Complete Partitioning Guide for IBM @server pSeries Servers*, SG24-7039
- ▶ *IBM @server Cluster 1600 Managed by PSSP 3.5: What's New*, SG24-6617
- ▶ *IBM @server pSeries 670 and pSeries 690 System Handbook*, SG24-7040
- ▶ *Linux Applications on pSeries*, SG24-6033
- ▶ *Linux Handbook: A Guide to IBM Linux Solutions and Resources*, SG24-7000
- ▶ *The POWER4 Processor Introduction and Tuning Guide*, SG24-7041
- ▶ *Practical Guide for SAN with pSeries*, SG24-6050
- ▶ *RS/6000 ATM Cookbook*, SG24-5525
- ▶ *RS/6000 SP and Clustered IBM @server pSeries Systems Handbook*, SG24-5596

AIX product manuals

The following types of documentation are located through the Internet at the following URL:

<http://www.ibm.com/servers/eserver/pseries/library>

- ▶ User guides
- ▶ System management guides
- ▶ Application programmer guides

- ▶ All commands reference volumes
- ▶ Files reference
- ▶ Technical reference volumes used by application programmers

Other publications

These publications are also relevant as further information sources:

- ▶ *7014 Series Model T00 and T42 Rack Installation and Service Guide*, SA38-0577
- ▶ *Hardware Management Console for pSeries Installation and Operations Guide*, SA38-0590
- ▶ *Hardware Management Console for pSeries Maintenance Guide*, SA38-0603
- ▶ *IBM PCI Cryptographic Accelerator Installation and Using Guide*, SA23-1254
- ▶ *PCI Cryptographic Coprocessor Installation and Using Guide*, SA23-1235
- ▶ *RS/6000 IBM @server pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*, SA38-0516
- ▶ *RS/6000 and IBM @server pSeries Adapter Placement Reference for AIX*, SA38-0538
- ▶ *RS/6000 and IBM @server pSeries Site and Hardware Planning Information*, SA38-0508
- ▶ *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280
- ▶ *System Unit Safety Information*, SA23-2652

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ *7040 System Planning Guide*, found at:
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm
- ▶ *AIX 5L Version 5.1 General Programming Concepts: Writing and Debugging Programs*, found at:
http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/genprog/genprogctfrm.htm

- ▶ *IBM @server pSeries 690 Configuring for Performance* whitepaper, found at:
http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_config.html
- ▶ *IBM @server pSeries and IBM RS/6000 Performance Report*, found at:
http://www.ibm.com/servers/eserver/pseries/hardware/system_perf.pdf
- ▶ *IBM @server pSeries Linux on pSeries Facts and Features* whitepaper, found at:
http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf
- ▶ *Linux for IBM @server pSeries: An overview for customers* whitepaper, found at:
http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html
- ▶ *POWER4 System Microarchitecture* whitepaper, found at:
<http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/power4.html>
- ▶ *pSeries 650 Installation Guide and Service Guide*, found at:
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm
- ▶ Adobe Reader main page
<http://www.adobe.com/products/acrobat/readermain.html>
- ▶ AIX 5L operating system description
<http://www.ibm.com/servers/aix/os/52desc.html>
- ▶ AIX Linux toolbox FTP site
<ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox>
- ▶ AIX Software page
<http://www.ibm.com/servers/aix/products>
- ▶ AIX Support For Large Page resource page
http://www.ibm.com/servers/aix/whitepapers/large_page.html
- ▶ AIX toolbox for Linux applications
<http://www.ibm.com/servers/aix/products/aixos/linux>
- ▶ AIX toolbox for Linux applications download page
<http://www.ibm.com/servers/aix/products/aixos/linux/download.html>
- ▶ alphaWorks Self Voicing Kit home page
<http://www.alphaworks.ibm.com/tech/svk>

- ▶ Apache Software Foundation home page
<http://www.apache.org>
- ▶ APC power device selector
<http://sturgeon.apcc.com/sizing.nsf/search>
- ▶ APC products
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=165
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=17
http://www.apcc.com/products/runtime_by_family.cfm?upsfamily=23
- ▶ Asian font packs for Adobe Reader
<http://www.adobe.com/products/acrobat/acrrasianfontpack.html>
- ▶ Autoupdate Perl script
<http://freshmeat.net/projects/autoupdate>
- ▶ Available code levels for SSA devices
<http://www.storage.ibm.com/hardsoft/products/ssa/rs6k/index.html>
- ▶ Avocent home page
<http://www.avocent.com/web/en.nsf>
- ▶ Beowulf Linux support information
<http://www.beowulf.org>
- ▶ Bull's Large Open Source Software Archive for AIX
<http://www.bullfreeware.com>
- ▶ Capacity Upgrade on Demand
<http://www.ibm.com/servers/eserver/pseries/cuod>
- ▶ A description of the pSeries cryptography products
<http://www.ibm.com/security/cryptocards>
- ▶ Facts and Features Reference
<http://www.ibm.com/servers/eserver/pseries/hardware/factsfeatures.html>
http://www.ibm.com/servers/eserver/pseries/hardware/workstations/facts_features.html
- ▶ Fast Connect for POWER Version 3.1 home page
http://www.ibm.com/servers/aix/products/ibmsw/system_man/fastconn.html
- ▶ FastT200 home page
<http://www.ibm.com/storage/FASTT200>
- ▶ FastT600 home page
<http://www.ibm.com/storage/fast600>

- ▶ FastT700 home page
<http://www.ibm.com/storage/fast700>
- ▶ FastT900 home page
<http://www.ibm.com/storage/fast900>
- ▶ The GNU Project home page
<http://www.gnu.org/gnu/the-gnu-project.html>
- ▶ GNU General Public License
<http://www.fsf.org/copyleft/gpl.html>
- ▶ IBM 2104 interoperability page
<http://www.storage.ibm.com/disk/expplus/supserver.htm>
- ▶ IBM AIX 5L
<http://www.ibm.com/servers/aix>
- ▶ IBM developer kits for AIX, Java technology edition
<http://www.ibm.com/developerworks/java/jdk/aix/index.html>
- ▶ IBM developer resources
<http://www.developer.ibm.com/spc/index.html>
- ▶ IBM Directory Portfolio
<http://www.Ibm.com/software/network/directory/server>
<http://www.ibm.com/software/network/directory/library>
- ▶ IBM disk storage products
<http://www.storage.ibm.com/hardsoft/disk/products.htm>
- ▶ IBM @server Cluster 1600
<http://www.ibm.com/servers/eservers/clusters/hardware/1600.html>
- ▶ IBM @server Cluster software resource page
<http://www.ibm.com/servers/eserver/clusters/software>
- ▶ IBM @server Support page
<http://techsupport.services.ibm.com/server/library>
- ▶ IBM HTTP server
<http://www.ibm.com/software/webservers/httpservers>
- ▶ IBM Linux Technology Center
<http://www.ibm.com/linux/ltc>

- ▶ IBM p690 RAS discussion
http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_ras.html
- ▶ IBM storage products
<http://www.ibm.com/storage/disk>
- ▶ IBM UNIX Servers: Hardware Documentation
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm
- ▶ IP Storage (ips) Charter
<http://www.ietf.org/html.charters/ips-charter.html>
<http://ietf.org/html.charters/ips-charter.html>
- ▶ iSCSI storage home page
<http://www.iscsistorage.com>
- ▶ Java Authentication and Authorization Service (JAAS)
<http://www.java.sun.com/products/jaas>
- ▶ Java Cryptography Extension (JCE)
<http://www.java.sun.com/products/jce>
<http://java.sun.com/products/jce>
- ▶ Java Secure Socket Extension (JSSE)
<http://www.java.sun.com/products/jsse>
<http://java.sun.com/products/jsse>
- ▶ Kerberos: The Network Authentication Protocol
<http://web.mit.edu/kerberos/www>
- ▶ Linear Tape-Open home page
<http://www.ibm.com/storage/lto>
- ▶ Linux at IBM resource page
<http://www.ibm.com/linux>
- ▶ Linux Kernel Archives
<http://www.kernel.org>
- ▶ Linux on pSeries home page
<http://www.ibm.com/servers/eserver/pseries/linux>
- ▶ Linux support line
http://www.ibm.com/services/e-business/linux_8.html
- ▶ Logical partitioning resource page
<http://www.ibm.com/servers/eserver/pseries/lpar/wp.html>

- ▶ MGE UPS Systems
<http://extranet.mgeups.com/accounts/ibm>
- ▶ Microcode Discovery Software
<http://techsupport.services.ibm.com/server/aix.invsoutMDS>
- ▶ Microcode download page for tape products
<http://techsupport.services.ibm.com/server/mdownload/tapehdr.html>
- ▶ Modular I/O
http://www.research.ibm.com/actc/Opt_Lib/mio/mio_doc.htm
- ▶ Myrinet on IBM Bladecenter information
<http://www.myricom.com>
- ▶ Netscape Communicator home page
<http://wp.netscape.com/browsers/4/index.html>
- ▶ NotesBench Consortium
<http://www.notesbench.org>
- ▶ Open Source Development Labs
<http://www.osdl.org>
- ▶ Open Source WAP and SMS gateway
<http://www.kannel.3glab.org>
- ▶ OpenSSH home page
<http://www.openssh.org>
- ▶ Parallel System Support Programs (PSSP) documentation
http://www.ibm.com/servers/eserver/pseries/library/sp_books/pssp.html
- ▶ Partnerworld developer resources
<http://www.developer.ibm.com/tech>
- ▶ PKCS #11 - Cryptographic Token Interface Standard
<http://www.rsasecurity.com/rsalabs/pkcs/pkcs-11>
- ▶ Powerware home page
<http://www.powerware.com>
- ▶ Powerware 9910 solutions
<http://www.oem.powerware.com/ibm-ups/9910solutions.html>
- ▶ Powerware pSeries solutions
<http://www.oem.powerware.com/ibm-ups/pdf/rsdream.pdf>

- ▶ pSeries Express Configurations
<http://www.ibm.com/servers/eserver/pseries/hardware/express.html>
- ▶ pSeries hardware documentation
http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware.htm
- ▶ pSeries microcode updates
<http://www.austin.ibm.com/support/micro>
- ▶ pSeries microcode updates download page
<http://techsupport.services.ibm.com/server/mdownload/download.html>
- ▶ pSeries and RS/6000 Microcode Updates
<http://techsupport.services.ibm.com/server/mdownload>
- ▶ pSeries and RS/6000 microcode updates download center
<http://www.rs6000.ibm.com/support/micro/download.html>
- ▶ Public Domain Software Library for AIX
<http://aixpdslib.seas.ucla.edu/aixpdslib.html>
- ▶ Red Hat Enterprise Server Linux
<http://www.redhat.com/software/eserver/pseries>
- ▶ Request For Comments home page
<http://www.ietf.org/rfc.html>
- ▶ RSA Public-Key Cryptography Standards
<http://www.rsasecurity.com/rsalabs/pkcs/index.html>
- ▶ SAN Switches for Core-to-Edge SAN solutions resource library
<http://www.storage.ibm.com/ibmsan/products/2109/library.html#support>
- ▶ SCSI and RAID controllers
<http://www.pc.ibm.com/us/compat/controllers/matrix>
- ▶ Sendmail home page
<http://www.sendmail.org>
- ▶ SP documentation
http://www.rs6000.ibm.com/resource/aix_resource/sp_books
- ▶ Standard Performance Evaluation Corporation
<http://www.spec.org>
- ▶ SuSE Linux Enterprise Server for pSeries
http://www.suse.com/us/business/products/server/sles/i_pseries.html

- ▶ Technical resources for storage products
<http://searchstorage.techtarget.com/whitepapers>
- ▶ Tivoli software home page
<http://www.tivoli.com/products/index/storage-mgr>
- ▶ Transaction Processing Performance Council home page
<http://www.tpc.org>
- ▶ TurboLinux for pSeries
<http://www.turbolinux.com/products/eserver/#pSeries>
- ▶ United Linux home page
<http://www.unitedlinux.com>
- ▶ UPS application sizing guide
<http://extranet.mgeups.com/accounts/ibm/doc/xlsizing.pdf>
- ▶ VERITAS Products A to Z
<http://www.veritas.com/products/listing/ProductListing.jhtml>
- ▶ VPD Capture Service
<http://techsupport.services.ibm.com/server/aix.invsoutVPD>
- ▶ Web-based System Manager home page
<http://www.ibm.com/servers/aix/wsm>
- ▶ Web browsers for AIX
<http://www.ibm.com/servers/aix/browsers/index.html>
- ▶ WebSphere software
<http://www.ibm.com/software/webservers/appserv>
- ▶ X.509 Public Key Infrastructure Certificate Management Protocols
<http://www.ietf.org/rfc/rfc2510.txt>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

/proc file system 534

Numerics

- 0920 AIX 5L Version 5.1 Bonus pack 589
- 0921 AIX 5L Version 5.1 Expansion pack 589
- 0948 AIX 5L Version 5.2 Bonus pack 589
- 0949 AIX 5L Version 5.2 Expansion pack 589
- 128-port asynchronous adapter
 - Remote Asynchronous Node (RAN) 391
- 15-pin D-shell connector, graphics accelerator 417
- 1722-60U 473
- 1742 TotalStorage FAStT900 Storage Server 465
- 1742-1RU 470
- 2104 458
- 2104,supported adapters 461
- 2104-DU3 458
- 2104-TU3 458
- 2105 Model 800 461
- 2105 Model 800 with expansion frame 464
- 2105 Model 800,advanced copy functions for business continuance 463
- 2109-F16 515
- 2109-F32 515
- 2109-M12 516
- 2109-S08 513
- 2109-S16 513
- 2498 PCI 4-Channel Ultra3 SCSI RAID Adapter 375
- 2623 4.7 GB SCSI-2 DVD-RAM Drive 434
- 2624 434
- 2624 32X (MAX) SCSI-2 CD-ROM Drive 434
- 2628 48X (Max) SCSI-2 Internal Auto-docking CD-ROM Drive 435
- 2629 4.7 GB SCSI-2 Auto-docking DVD-RAM Drive 436
- 2633 437
- 2634 16X/48X(max) IDE DVD-ROM Drive 437
- 2635 16X/48X(max) SCSI Auto-docking DVD-ROM Drive 438
- 2640 IDE Slimline DVD-ROM Drive 439
- 2732 IBM Short-wave Serial HIPPI PCI Adapter 386
- 2751 S/390 ESCON Channel Adapter 387
- 2826 GXT4000P Graphics Accelerator 419
- 2827 GXT6000P Graphics Accelerator 423
- 2842 GXT4500P Graphics Accelerator 421
- 2843 GXT6500P Graphics Accelerator 425
- 2848 GXT135P Graphics Accelerator 417
- 2944 128-Port Asynchronous Adapter EIA-232 391
- 2946 Turboways 622 Mbps PCI MMF ATM Adapter 399
- 2947 ARTIC960Hx 4-Port Selectable Adapter 393
- 2962 2-Port Multiprotocol X.25 Adapter 397
- 2969 Gigabit Ethernet - SX Adapter 400
- 2975 10/100/1000 Base-T Ethernet PCI Adapter 401
- 32-bit 3
- 3494 487
 - components 488
 - control unit frame 488
 - drive unit frame 489
 - hardware requirements 490
 - highlights 488
 - number of tape drives 490
 - operating systems 490–491
 - storage unit frame 489
- 3494-D12
 - capacity 490
- 3494-L12
 - capacity 490
 - cartridge capacity 490
 - number of drive unit frames 490
- 3534-F08 513
- 3542-1RU 476
- 3542-1RX 476
- 3542-2RU 476
- 3542-2RX 476
- 3580 492
- 3580 Ultrium 1 drive
 - capacity per cartridge 493
 - dimensions 494
 - interface 493
 - media 494
 - operating systems 493
 - warranty 494
 - weight 494

- 3580 Ultrium 2 drive
 - capacity per cartridge 494
 - dimensions 495
 - interface 494
 - media 495
 - operating systems 495
 - warranty 495
 - weight 495
- 3581 496
 - capacity 497
 - dimensions 498
 - highlights 496
 - interface 497
 - media 498
 - number of tape cartridges 497
 - operating systems 498
 - tape drive type 497
 - warranty 498
 - weight 498
- 3582 498
 - capacity per cartridge 499
 - dimensions 500
 - highlights 499
 - interface 499
 - media 500
 - number of tape cartridges 499
 - number of tape drives 499
 - operating systems 500
 - warranty 500
 - weight 500
- 3583 501
 - capacity 504
 - compatibility 503
 - configuration notes 503
 - dimensions 506
 - highlights 502
 - interface 505
 - media 506
 - number of tape drives 504
 - operating systems 506
 - tape drive type 504
 - warranty 507
 - weight 506
- 3584 507
 - capacity 510
 - components 507
 - configuration notes 509
 - dimensions 511
 - highlights 508
 - interface 510
 - media 511
 - model conversion 508
 - number of drives 509
 - number of frames 509
 - number of logical libraries 509
 - number of tape cartridges 509
 - operating systems 510
 - tape drive type 509
 - warranty 511
 - weight 511
- 3590 485
 - model description 486
- 4 mm Internal Tape Drive 441
- 4012 SP Switch2 Adapter 361
- 4023 SP Switch MX2 Adapter 361
- 4025 SP Switch2 Adapter 361
- 4026 SP Switch2 MX2 Adapter 361
- 4032 SP Switch2 Interposer 362
- 4217 6091 Cable 6ft 15-D-to-5BNC 428
- 4237 Cable P201 15-D-to-13W3 - DDC/1010 428
- 4238 CABLE, 15P D-SHELL TO 13W3, W/DDC, FOR P70/200 428
- 4239 15-D ADAPTER TO 3BNC ID CABLE 428
- 4951 4-Port 10/100 Base-Tx Ethernet PCI Adapter 402
- 4951 IBM 4-Port 10/100 Ethernet Adapter 404
- 4953 IBM 64bit/66MHz PCI ATM 155 UTP Adapter 399
- 4957 IBM 64bit/66MHz PCI ATM 155 MMF Adapter 399
- 4959 4/16 Mbps Token-Ring Adapter 403
- 4960 IBM e-business Cryptographic Accelerator 411
- 4961 IBM Universal 4-Port 10/100 Ethernet Adapter 405
- 4962 10/100 Mbps Ethernet PCI Adapter II 406
- 4963 PCI Cryptographic Coprocessor-FIPS-4 412
- 4-way 1.1GHz POWER4 218
- 4-way 1.3GHz POWER4 HPC 218
- 5700 IBM Gigabit Ethernet-SX PCI-X Adapter 407
- 5701 IBM 10/100/1000 Base-TX Ethernet PCI-X Adapter 407
- 5706 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter 408
- 5707 IBM 2-Port Gigabit Ethernet-SX PCI-X Adapter 409
- 6120 IBM 80/160 GB Internal Tape Drive 440
- 6131 60/150 GB 16-bit 8mm Internal Auto-docking

Tape Drive 440
 6134 60/150 GB 16-bit 8mm Internal Tape Drive 441
 6158 20/40GB 4mm Internal Tape Drive 441
 6169 IBM 80/160 GB Internal Auto-docking Tape Drive 442
 6185 20/40 GB 4mm Internal Auto-docking Tape Drive 443
 6187 Power Controller 352
 6188 Power Distribution Assembly 352
 6200 Integrated Battery Backup Primary 352
 6201 Integrated Battery Backup Redundant 352
 6203 PCI Dual Channel Ultra3 SCSI Adapter 377
 6204 PCI Universal Differential Ultra SCSI Adapter 377
 6205 PCI Dual Channel Ultra2 SCSI Adapter 378
 6206 PCI Single-Ended Ultra SCSI Adapter 379, 918
 6206 PCI single-ended Ultra SCSI Adapter 919, 924
 6206 PCI Single-Ended Ultra-SCSI Adapter 918
 6207 PCI Differential Ultra SCSI Adapter 381
 6228 Gigabit Fibre Channel Adapter for 64-bit PCI Bus 386
 6230 Advanced SerialRAID Plus Adapter 382
 6239 2 Gigabit Fibre Channel PCI-X Adapter 410
 6310 ARTIC960RxD Quad Digital Trunk Adapter 394
 6312 Quad Digital Trunk Telephony PCI Adapter 396
 64-bit 3
 64-bit kernel 525
 7006 IBM RS/6000, system life time information 956
 7007 IBM RS/6000, system life time information 956
 7008 IBM RS/6000, system life time information 956
 7009 IBM RS/6000, system life time information 957
 7010 IBM RS/6000, system life time information 957
 7011 IBM RS/6000, system life time information 957
 7012 IBM RS/6000, system life time information 958
 7013 IBM RS/6000, system life time information 959
 7014 IBM RS/6000, system life time information 960
 7014-T00 310
 7014-T42 317
 7015 IBM RS/6000, system life time information 960
 7017 IBM RS/6000, system life time information 961
 7024 IBM RS/6000 Server, system life time information 962
 7025 IBM RS/6000 Server, system life time information 962
 7025 Model 6F0,adapter placement guidelines 789
 7025 Model 6F1,adapter placement guidelines 789
 7025 pSeries620 Models 6F0 and 6F1,adapter placement guidelines 789
 7026 IBM RS/6000 Server, system life time information 963
 7026 Model B80,adapter placement guidelines 798
 7026 pSeries614 Model B80,adapter placement guidelines 798
 7027 IBM High Capacity Drawer, system life time information 963
 7028 IBM pSeries 610 Server, system life time information 964
 7028 Model 6C1,adapter placement guidelines 803
 7028 Model 6C4 (4-Slot PCI riser), adapter placement guidelines 809
 7028 Model 6C4 (6-Slot PCI riser), adapter placement guidelines 815
 7028 Model 6E1,adapter placement guidelines 803
 7028 Model 6E4 (4-Slot PCI riser), adapter placement guidelines 809
 7028 Model 6E4 (6-Slot PCI riser), adapter placement guidelines 815
 7029 Model 6C3,adapter placement guidelines 820
 7029 Model 6E3,adapter placement guidelines 820
 7030 IBM RS/6000, system life time information 964
 7038 Model 6M2,adapter placement guidelines 845
 7038-6M2 IBM pSeries 650 Server Model 6M2, system life time information 964
 7039 Model 651
 7040-61D I/O Drawer 350
 7040-W42 System Rack 352
 7042-W42 347
 AIX 348
 Cluster 1600 346
 disk and boot devices 351

- feature codes 353
- Hardware Management Console 347
- Memory 345
- POWER4+ 344
- PSSP 346
 - rules for racking p655 components 352
 - SP Switch2 PCI-X Attachment Adapter 349
- 7039 Model 651, adapter placement guidelines 850
- 7039-651 IBM pSeries 655 Server Model 651, system life time information 965
- 7040 IBM pSeries 670/690, system life time information 965
- 7040 Model 671, adapter placement guidelines 824
- 7040 Model 681, adapter placement guidelines 835
- 7040 pSeries 670 215
- 7040 pSeries 670 Model 671 242
- 7040 pSeries 670/690, features 295
- 7040 pSeries 690 215
- 7040 pSeries 690 Model 681 266
- 7040 pSeries technical overview 217
- 7040, primary I/O book 230
- 7040, secondary I/O books 230–231
- 7040-61D 232
- 7040-61R 237
- 7040-671 242
- 7040-671/681, CEC 219
- 7040-media drawer 236
- 7040-RAS features 241
 - fault recovery 242
 - predictive functions 242
 - redundancy in components 242
 - serviceability features 242
- 7043 IBM RS/6000 43P, system life time information 965
- 7043 Model 150, adapter placement guidelines 859
- 7044 IBM RS/6000, system life time information 966
- 7044 Model 170, adapter placement guidelines 867
- 7044 Model 270, adapter placement guidelines 876
- 7045-SW4 High Performance Switch
 - High performance switch (7045-SW4) 362
- 7046 Model B50, adapter placement guidelines 884
- 7046-B50 IBM RS/6000 Server, system life time information 966
- 7133 -D40 455
- 7133 -T40 455
- 7202-900 IBM RS/6000 Expansion Rack Model 900, system life time information 966
- 7205 480
 - capacity 481
 - data transfer rate 481
 - dimensions 481
 - interface 481
 - media 481
 - operating systems 481
 - warranty 481
 - weight 481
- 7206 481
- 7206-220
 - capacity 482
 - data transfer rate 482
 - dimensions 482
 - interface 482
 - media 482
 - operating systems 482
 - warranty 482
- 7206-VX2
 - capacity 483
 - data transfer rate 483
 - dimensions 483
 - interface 483
 - media 483
 - operating systems 483
 - warranty 483
- 7207 484
 - capacity 484
 - data transfer rate 484
 - interface 484
 - media 485
 - operating systems 485
 - warranty 485
- 7236-001 IBM RS/6000 Media Streamer Model 001, system life time information 967
- 7248 IBM RS/6000 43P Series, system life time information 967
- 7311 IBM rack-mounted drawers, system life time information 967
- 7311 Model D10, adapter placement guidelines 887
- 7311 Model D20, adapter placement guidelines 891
- 7311-D10 206
 - features 209
 - hardware minimum requirements 208
 - highlights 207
 - I/O drawer attachment 208
 - maximum number 207
 - PCI slots 207
- 7311-D20
 - configuration notes 204

disk drive bays 203
 features 209
 hardware minimum requirements 204
 highlights 203
 internal disk drives 203
 maximum number 204
 PCI slots 203
 remote I/O cables 204–205
 RIO ports 204
 RIO-2 ports 204
 SPCN cables 204, 206
 7315-C01 IBM Hardware Management Console,
 system life time information 968
 7316 IBM rack-mounted console, system life time in-
 formation 968
 7316-TF2 325
 7317 IBM RS/6000 telecom server and external scsi
 disk, system life time information 968
 7318 IBM serial communications network server,
 system life time information 969
 7319 IBM RS/6000 fibre channel switch and adapt-
 er, system life time information 969
 8 mm Internal Tape Drive 441
 8 way 1.3GHz POWER4 Turbo 218
 8396 IBM RS/6000 SP System Attachment Adapter
 388
 8397 SP Switch2 PCI Attachment Adapter 389
 8398 SP Switch2 PCI-X Attachment Adapter 389
 8677 line Cord, WT, 50A/240V, 30A/480V, 8AWG,
 14ft, No Plug 909
 8678 line Cord, 60A, 240V, 6AWG, 14ft, US, Chica-
 go IEC309 Plug 909
 8680 line Cord, US, Canada, Japan, 480V,
 30A, 10AWG, 14ft, IEC309 Plug 909
 8681 Line Cord, US, Canada, Japan, 240V, 60A,
 6AWG, 6ft, Chicago IEC309 Plug 909
 8682 line Cord, US, Chicago, 480V, 30A, 10AWG,
 6ft, Chicago, IEC309 Plug 909
 8683 line Cord, US, Canada, Japan, 280V, 60A,
 8AWG, 14ft, IEC309 Plug, Single Phase 909
 8684 line Cord, 8AWG, 6ft, Chicago, IEC309 Plug,
 Single Phase, 60A, 200-280V 909
 8685 line Cord, WT, 40A, 200-415V, 8AWG, 14ft,
 No Plug, Single Phase 909
 8686 line Cord, US, Canada, Japan, 200-240V,
 6AWG, 14ft, IEC309 100A Plug 909
 8687 line Cord, US, Chicago, 200-240V, 6AWG, 6ft,
 IEC309 100A Plug 909
 8688 line Cord, US, Canada, Japan, 200-240V,
 6AWG/Type W, 14ft, IEC309 60A Plug 909
 8689 line Cord 909
 8-way 1.1GHz POWER4 218
 9076 Model 555 359
 9076 Model 556 359
 9076 Model 557 360
 9076 Model 558 360
 9076 RS/6000 SP Systems, adapter placement
 guidelines 897
 9112 Model 265, adapter placement guidelines 897
 9800 power Cord - United States/Canada, L6-30P,
 250V, single phase, 14 ft., 30A 908
 9802 power cord - Brazil (125V, 15A) 908
 9820 power cord - Belgium, Finland, France (250V,
 16A) 908
 9821 power cord - Denmark (250V, 10A) 908
 9825 power cord - United Kingdom (250V, 13A)
 908
 9827 power cord - Israel (250V, 6-16A) 908
 9828 power cord - Switzerland (250V, 10A) 908
 9829 power cord - South Africa/Pakistan (250V,
 16A) 908
 9830 power cord - Italy (250V, 10A and 16A) 908
 9831 power cord - Argentina (250V, 10A) 908
 9833 power cord - Thailand (250V, 15A) 908
 9834 power cord - Uruguay (250V, 10A) 908
 9900 power cord (4M) - United States/Canada 908
 9901 power cord (4M) - Belgium, Finland, France
 (250V, 16A) 908
 9902 power cord (4M) - Denmark (250V, 10A) 908
 9903 power cord (4M) - U.K. and Others (250V,
 13A) 908
 9904 power cord (4M) - Israel (250V, 6-16A) 908
 9905 Power cord (4M) - Switzerland (250V, 10A)
 908
 9906 power cord (4M) - India, Pakistan, S. Africa
 (250V, 16A) 908
 9907 power cord (4M) - Italy (250V, 10A and 16A)
 908
 9908 power cord (4M) - Australia/New Zealand/Ar-
 gentina (250V, 10A) 908
 9909 power cord (4M) - Thailand (250V, 15A) 908
 9910-A08 331
 9910-A10 331
 9910-A13 332
 9910-A30 333
 9910-Axx 330
 9910-M10 339
 9910-M11 339

9910-M15 340
 9910-M30 341
 9910-Mxx 338
 9910-P07 335
 9910-P18 335
 9910-P22 336
 9910-P33 337
 9910-P60 338
 9910-Pxx 333
 9911 power cord (4M) - All (Standard rack power cord) 908

A

ABI, application binary interface 526
 AC power distribution bus 313
 adapter firmware download 971
 adapter placement guidelines 785
 7025 pSeries 620 Models 6F0 and 6F1 789
 7026 pSeries 640 Model B80 798
 7028 pSeries 610 Models 6C1 and 6E1 803
 7028 pSeries 630 Models 6C4 and 6E4 (4-Slot PCI Riser) 809
 7028 pSeries 630 Models 6C4 and 6E4 (6-Slot PCI Riser) 815
 7028 pSeries630 Models 6C4 and 6E4
 Logical partition (LPAR) considerations 809
 7029 pSeries 615 Models 6C3 and 6E3 820
 7029 pSeries615 Models 6C3 and 6E3
 Logical partition (LPAR) considerations 820
 7038 pSeries 650 Model 6M2 845
 Logical partition (LPAR) considerations 846
 7039 pSeries 655 Model 651 850
 Logical partition (LPAR) considerations 850
 7040 pSeries 670 Model 671 drawers 61D 824
 7040 pSeries 690 Model 681 drawer 61D 835
 7040 pSeries670 Model 671 Drawers 61D
 Logical partition (LPAR) considerations 827
 7040 pSeries690 Model 681 drawer 61D
 Logical partition (LPAR) considerations 837
 7043 Model 150 859
 7043 Model 150 Multiple Graphics Adapter Placement Guide 864
 7044 Model 170 867
 7044 Model 170 Multiple Graphics Adapter Placement Guide 874
 7044 Model 270 876
 7044 Model 270 Multiple Graphics Adapter Placement Guide 882

7046 Model B50 884
 7311 Model D10 887
 Logical partition (LPAR) considerations 887
 7311 Model D20 891
 Logical partition (LPAR) considerations 892
 9076 RS/6000 SP Systems 897
 9112 Model 265 897
 adapters, communications
 8-port Asynchronous Adapter 390
 Adobe Acrobat Reader 5.06 582
 affinity logical partitions 710
 Air Moving Device,7040 240
 AIX
 AIX 5L Version 5.1 522
 /proc file system 534
 64-bit kernel 525
 ABI, application binary interface 526
 AIX LDAP security audit plug-in 552
 AIX system security 550
 ANSI terminal support 535
 argument list limit 535
 automatic dump analysis 541
 Cachefs 538
 Certificate Management Protocol (CMP) 543
 Code Page 943 531
 comparison between JFS2 and JFS 539
 concurrent group 534
 core files 541
 dedicated dump device 541
 Development and Performance Tools 542
 DMAPI, Data Management API 537
 dump reliability improvements 541
 Dynamic Feedback Protocol (DFP) 546
 error log retrieval API 540
 error log scalability 539
 Euro currency symbol support 531
 event notification support 551
 Explicit Congestion Notification (ECN) 546
 GBK locale 533
 hardware requirements 522
 highlights 525
 hot spare disk support 538
 hot spot management 538
 IBM AIX Developer Kit, JAVA2 Technology 542
 IBM Self-Voicing Kit (SVK) 529
 ICU, Internationalized Classes for Unicode 532

- International Language Support 529
- Inventory Scout 545
- IP fragmentation 549
- IP Key Encryption Security 554
- IP multi-routing and multiple gateway 546
- ISO Standard ISO8859-15 532
- Java Cryptography Extension (JCE) 543
- Java Secure Sockets Extension (JSSE) 543
- Java security 542
- JFS2 538
- Kerberos Version 5 Support 552
- Kernel Debugger (KDB) 541
- large program support 535
- lg_dumplv 541
- Linux affinity 545
- logical track group (LTG) 537
- malloc enhancements 533
- MBCS, Multi-Byte Character Set Support 529
- Mirror Write Consistency (MWC) 537
- multithreaded autoFS 534
- Native Kerberos Version 5 KDC Server/Client Support 553
- network 546
- network interface backup mode 547
- networking security 554
- new locales 532
- NFS statd multithreading 533
- packet capture library 548
- paging space 533
- PASSTHRU Support 529
- PKCS support 554
- Pluggable Authentication Mechanism (PAM) 553
- Public Key Cryptography Standards (PKCS) 543
- Quality of Service (QoS) 549
- RAS, Reliability, Availability, and Serviceability 539
- release and withdrawal date 525
- resource management 535
- RSCT, Reliable Scalable Cluster Technology 526
- SBCS, Single Byte Character Set Support 529
- SecureWay Directory Server 524
- Security Audit Support 551
- Sendmail Version 8.11.0 548
- software requirements 524
- storage management 537
- SVR4 printing subsystem 534
- system hang detection 539
- system management 525
- TCP splicing 547
- Token Ring support for MPOA 549
- transaction support 552
- Unicode support 530
- Virtual IP Address (VIPA) 547
- Virtual Local Area Network (VLAN) 549
- Web serving enhancements 546
- Workload Manager 535
- WSM application support 527
- WSM, Web-based System Manager 526
- AIX 5L Version 5.2 555
 - advanced RAID support 568
 - AIX system security 573
 - AIX tuning framework 571
 - alternate disk installation 564
 - boot LED displays 570
 - Cluster Systems Management (CSM) 563
 - Common HBA API support 568
 - cryptographically secure random number 574
 - development and performance tools 570
 - Diffie-Hellman group 5 575
 - Dynamic Capacity Upgrade on Demand (DCUoD) 565
 - dynamic CPU guard 566
 - Dynamic LPAR 565
 - EtherChannel backup adapter 572
 - file descriptor limit 567
 - generic data management tunnel 574
 - hardware requirement 555
 - highlights 563
 - HMC recovery software Version 1.3 562
 - IBM AIX Developer Kit, JAVA2 Technology 570
 - IBM Directory 562
 - IBM Directory Server 574
 - JFS2 snapshot image 569
 - Linux affinity 571
 - Mobile IPv6 572
 - multipath I/O 570
 - network 571
 - networking security 574
 - PAM enhancements 573
 - performance analysis tools 570
 - Perl 5.8 571

- Public Key Infrastructure enhancements 573
- Quality of Service (QoS) 573
- RAS, Reliability, Availability, and Serviceability 569
- release and withdrawal date 563
- resource management 565
- SNMPv3 571
- software requirement 561
- storage management 567
- system dump facility enhancements 569
- system hang detection 570
- system management 563
- unaligned I/O support in LVM 568
- Universal Disk Format (UDF) 567
- withdrawn device support 561
- withdrawn ISA adapter support 558
- withdrawn MCA RS/6000 support listing 559
- withdrawn MCA-based SP nodes support 560
- withdrawn PCI adapter support 556
- withdrawn PCI RS/6000 support listing 558
- withdrawn PReP-specific ISA adapter support 557
- Workload Manager (WLM) 566
- Xprofiler analysis tool 570
- AIX 5L Bonus Pack 585
- AIX 5L Expansion Pack 584
- AIX 5L Expansion, Bonus, and Web Download pack 575
 - AIX iSCSI Initiator Version 1.0 584
 - AIX Toolbox for Linux application 581
 - Certificate Authentication Service Version 5.2.0.10 576
 - contents summary 584
 - cryptographic library Version 5.2 576
 - Data Encryption Standard (DES) library routines 578
 - development 580
 - e-business and e-business 579
 - encryption support for IBM Directory Version 4.1 577
 - encryption support for SecureWay Directory Version 3.2.2 577
 - Geodesic Systems Great Circle Version 6.0.1.5 581
 - how to order 588
 - IBM AIX Developer Kit, JAVA2 Technology 580
 - IBM Directory Server 579
 - IBM Global Security Toolkit Version 4.0 577
 - IBM Global Security Toolkit Version 5.0 577
 - IBM HTTP Server 579
 - IBM IP Security Version 5.1 and Version 5.2 578
 - IBM Network Authentication Service Version 1.3 576
 - IBM Tivoli Storage Manager Version 5.1.5 583
 - IBM Tivoli Storage Resource Manager Version 1.2 583
 - interoperability 582
 - Modular I/O (MIO) Library Version 2.3 581
 - Multimedia 582
 - Netscape Communicator 4.8 580
 - Open Secure Shell (OpenSSH) Version 3.6 578
 - PartnerWorld for developers 582
 - Reliable Scalable Cluster Technology Version 2.3 577
 - SNMPv3 Encryption Version 5.2 578
 - storage management 583
 - system management 576
 - VERITAS Foundation Suite 583
 - Web-based System Manager Security Version 5.1 and Version 5.2 578
 - WebSphere Application Server 579
 - Wireless Application Protocol (WAP) Gateway Version 1.1 580
- AIX 5L Version 5.1
 - performance analysis tools 543
- AIX 5L Web Download Pack 587
- AIX Fast Connect Version 3.1.2 582
- AIX iSCSI Initiator Version 1.0 584
- AIX system security 550
- AIX Toolbox for Linux application 581, 604
 - archiving applications 606
 - database application 606
 - development libraries 605
 - development tools 605
 - editing applications 606
 - graphical desktops 605
 - graphical games 606
 - internet applications 606
 - Open Source Software 605
 - programming languages 606
 - system applications 606
 - system shells 606
 - user interface for desktops 605
- ANSI terminal 535

- anti-aliasing, graphics accelerator 416
- Apache Server 1.3.12 562
- asynchronous adapters
 - 8-port Asynchronous Adapter 390
- Auto-docking Tape Drive 442
- automatic dump analysis 541
- Autonomic Computing 10
- Autonomic computing
 - Capacity Upgrade on Demand 16
 - e-business on demand 12
 - On demand operating environment 13
 - RAS 14
 - self-configuring 11
 - self-healing 11
 - self-optimization 11
 - self-protecting 12
- autonomic computing 725

B

- Backup File Format (BFF) 607
- Berkeley Packet Filter (BPF) 548
- boot adapters and IP address takeover 630
- BPA,7040 238
- BPC,7040 238
- BPD,7040 238
- BPR,7040 238
- building block 687
- Bulk Power Assembly,7040 238
- Bulk Power Controller,7040 238
- Bulk Power Distributor,7040 238
- Bulk Power Regulator,7040 238

C

- CA, Certificate Authority 573
- cables, SCSI (various types) 919, 924
- Cachefs 538
- capacitor book,7040 220
- Capacity BackUp 17
- Capacity Upgrade on Demand 16
- capacity upgrades,2105 Model 800 465
- cascading resource groups 633
- Cascading Without Fallback (CWOFF) 633
- CD/DVD devices 695
- CD-ROM 434
- CD-RW 434
- CEC 7040-671/681 219
 - L2 cache subsystem 224
 - L3 cache controller and directory 225

- memory subsystem for pSeries 670 228
- multichip modules 226
- POWER4 core 224
 - single core POWER4 processor feature 227
- CEC backplane,7040 221
- CEC front view 220
- CEC rear view 221
- Central Electronics Complex 219
- Central Electronics Complex (CEC 7040-671/681) 219
- Certificate Authentication Service Version 5.2.0.10 576
- Certificate Management Protocol (CMP) 543
- Certificate Revocation Lists (CRL) 554
- Certification Management Protocol (CMP) 580
- chipkill 599
- CIM 728
- CIM object manager 728
- CIU, Core Interface Unit 8
- CIU,7040 225
- Class I, graphics accelerators 416
- Class II, graphics accelerators 416
- Class III, graphics accelerators 416
- clock card,7040 221
- Cluster 1600 672
 - Cluster 1600 with CSM 681
 - expansibility 681
 - prerequisites 683
 - Cluster 1600 with PSSP 678
 - prerequisites 680
- Cluster Systems Management (CSM) 674
- CSM hardware control features 677
- CSM interoperability 677
- Dual Hardware Management Console (HMC) support 677
- enhancements 676
- General Parallel File System (GPFS) 674
- models and features 674
- Parallel System Support Program (PSSP) 673
- pSeries 630 Server 677
- pSeries 650 cluster support 677
- pSeries 655 servers 676
- pSeries 670 and 690 server 676
 - supported hardware 678
- cluster single point of control (C-SPOC) 644
- cluster systems 627
- Cluster Systems Management (CSM) 563, 674
- CMOS-8S3SOI 224
- CMOS-9SSOI 224

Code Page 943 531
 commands
 acctcom 536
 alstat 543
 bf 544
 bfrpt 544
 bindintcpu 566
 bindprocessor 566
 bosboot -L 543
 chvg 537
 clverify 645
 compare_report 565
 curt 570
 emstat 543
 errpt 540
 filemon 538, 544, 570
 fileplace 570
 gennames 544
 installp 564
 iostat 538, 544–545
 kdb 541
 lockstat 544
 locktrace 543–544
 mkvg 537
 netpmon 544
 nfso 571
 nimadm 564
 no 571
 perfstat 544
 pprof 544
 proc 570
 restore 545
 rmss 544
 sar 545
 schedtune 571
 shdaemon 539
 shrinkps 533
 smit hacmp 646
 smit vlan 550
 smitty compare_report 565
 smitty nimadm 564
 snap 541
 splat 570
 stem 544
 svmon 544
 swapoff 533
 syscalls 544
 tar 545
 topas 544

 tprof 544
 truss 543–544, 570
 varyonvg rootvg 570
 vmstat 544
 vmtune 571
 wlmmon 544
 Common HBA API support 568
 Common Information Model 728
 concurrent resource groups 633
 Concurrent Resource Manager (CRM) 639
 Congestion Window Reduced (CWR) 547
 control workstation 719
 copper interconnects 7
 CPU Guard 711
 CRM, Customer Relationship Management 1
 cryptographic library Version 5.2 576
 Customer Installation Matrix and Processor Groups 947

D

Data Encryption Standard (DES) library routines 578
 Data Management Application Programming Interface (DMAPI) 537
 DB2 Version 7.2 562
 DCA,7040 220
 DDC2B support, graphics accelerator 422
 Denial of Service across shared resources 696
 device firmware download 971
 Diffie-Hellman group 5 575
 direct memory access 699, 703
 disk firmware download 971
 disk systems summary 452
 display 428
 6091-19I 428
 6656-Hxx T560 428
 6658-Hxx T84H 429
 6659-Hxx T210 429
 9493-Axx T56A 428
 9494-xxx T860 428
 9497-Axx T86 428
 9511-Axx T54A/T54H/ T540 428
 9516-B03 428
 9519-Axx T85A 428
 P200 428
 P201 428
 P202 428
 P260 428

- P275 428
- P70 428
- P72 428
- P76 428
- P77 428
- P92 428
- POWERDP17 428
- POWERDP20 428
- distributed converter assembly,7040 220
- Distributed Management Task Force 728
- DLPAR 696, 725
 - overview 725
- DMAPI, Data Management API 537
- DMTF 728
- domains 687
- download firmware 971
- DSP, Digital Signal Processor 416
- DVD-RAM 434
- DVI connector, graphics accelerator 417
- DWA, direct window access 523
- Dynamic Automatic Reconfiguration Events (DARE) 643
- Dynamic Capacity Upgrade on Demand (DCUoD) 565
- dynamic CPU guard 566
- Dynamic Feedback Protocol (DFP) 546
- dynamic host configuration protocol (DHCP) 574
- Dynamic logical partitioning 696, 710, 725
- Dynamic LPAR (DLPAR) 565
- dynamic processor deallocation 711
- Dynamic Reconfiguration (DR) 566

E

- e-business on demand 12
- ECC memory 599
- ECC, Error Checking and Correction 15
- EEH, partitioning 703
- eight execution unit 224
- Electronic Service Agent 735
- enhanced error handling, partitioning 703
- Enhanced Scalability Concurrent Resource Manager (ESCRM) 639
- Enterprise storage server 2105 Model 800 461
 - advanced copy functions for business continuance 463
 - capacity upgrades 465
 - highlights 462
 - physical characteristics 464

- Enterprise storage server 2105 Model 800,advanced copy functions for business continuance 463

Entry systems

- 7026 Model B80 IBM pSeries 640 139
- 7026 Model B80 pSeries 640
 - configuration notes 145
 - disks, media, and boot devices 147
 - Graphics adapters 149
 - hot-plug options 149
 - memory 145
 - PCI slots and adapters 148
 - power cords 146
 - power supplies 146
 - processors 145
 - racks 147
- features 145
- highlights 140
- minimum and standard features 143
- product positioning 140
- RAS features 143
- system expansion 144
- technical overview 141
- 7028 Model 6C1 and 6E1
 - express configurations 44
- 7028 Model 6C1 and 6E1 pSeries 610 32
 - configuration notes 38
 - disks, media, and boot devices 40
 - graphics adapters 44
 - hot-plug options 44
 - memory 39
 - PCI slots and adapters 43
 - power cords 39
 - power supplies 39
 - processors 38
 - racks 40
- features 38, 138
- highlights 33
- minimum and standard features 36
- product positioning 32
- RAS features 36
- system expansion 38
- technical overview 34
- 7028 Model 6C4 and 6E4
 - express configurations 98
- 7028 Model 6C4 and 6E4 pSeries 630 77
 - configuration notes 89
 - disks, media, and boot devices 96
 - graphics adapters 98

- hot-plug options 98
- I/O drawer attachment 94
- memory 91
- PCI slots and adapters 97
- power cords 93
- power supplies 93
- processors 90
- racks 94
- features 89
- highlights 79
- minimum and standard features 87
- product positioning 78
- RAS features 86
- system expansion 88
- technical overview 80
- 7029 Model 6C3 and 6E3
 - express configurations 65
- 7029 Model 6C3 and 6E3 pSeries 615 46
 - configuration notes 58
 - disks, media, and boot devices 62
 - graphics adapters 64
 - hot-plug options 65
 - memory 60
 - PCI slots and adapters 64
 - power cords 61
 - power supplies 61
 - processors 59
 - racks 61
 - features 58
 - highlights 49
 - minimum and standard features 56
 - product positioning 48
 - RAS features 55
 - system expansion 57
 - technical overview 50
- 7043 Model 150 IBM RS/6000 43P 112
 - configuration notes 118
 - disks, media, and boot devices 119
 - graphics adapters 122
 - memory 118
 - PCI slots and adapters 121
 - processors 118
 - features 118
 - highlights 113
 - minimum and standard features 116
 - product positioning 113
 - RAS features 116
 - system expansion 117
 - technical overview 114
- 7044 Model 170 IBM RS/6000 44P 122
 - configuration notes 130
 - disks, media, and boot devices 132
 - graphics adapters 134
 - memory 131
 - PCI slots and adapters 134
 - processors 131
 - features 130
 - highlights 125
 - minimum and standard features 129
 - product positioning 124
 - RAS features 129
 - system expansion 130
 - technical overview 126
- 9112 Model 265 IBM IntelliStation 134
 - configuration notes 138
 - graphics adapters 138
 - processors 138
 - express configuration 139
 - highlights 136
 - minimum and standard features 136
 - product positioning 136
 - RAS features 136
 - system expansion 138
 - technical overview 136
- 9114 Model 275 IBM IntelliStation 109
 - configuration notes 110
 - audio support 111
 - graphic adapters 110
 - graphics adapters 110
 - PCI slots and adapters 110
 - processors 110
 - express configuration 111
 - general description 109
- EPOW, Early Power Off Warning 351
- ERP, Enterprise Resource Planning 1
- error log retrieval API 540
- ESS 461
- EtherChannel 547, 572
- EtherChannel backup adapter 572
- Euro currency symbol support 531
- Expandable Storage Plus 2104 458
 - flexible and scalable storage 459
 - high availability 459
 - highlights 458
 - physical dimensions 460
 - supported adapters 461
- Explicit Congestion Notification (ECN) 546
- express configurations

- Intellistation 275 111
 - p610 44
 - p615 65
 - p630 98
 - p650 195
- Extended Remote Copy (XRC), Enterprise storage server 2105 Model 800 464
- external device support table, all RS/6000's 931
- external disk storage 452
- external disk systems summary 452
- external storage architectures and devices 451
 - disk systems summary 452
 - Enterprise storage server 2105 Model 800 461
 - Expandable Storage Plus 2104 458
 - IBM 7133 Serial Disk System Advanced Models T40 and D40 455
 - IBM external disk storage 452
 - IBM TotalStorage FAStT200 Storage Server 476
 - IBM TotalStorage FAStT600 Storage Server 473
 - IBM TotalStorage FAStT700 Storage Server 470
 - IBM TotalStorage FAStT900 Storage Server 465

F

- Fabric Controller,7040 225
- facts and features reference 19
 - server I/O attachments 26
 - software support 29
 - standard warranty 28
 - system unit details 25
- Fast reboot in a partitioned environment 709
- Fast Response Cache Architecture (FRCA) 546
- FAStT200 Storage Server 476
 - description 477
 - hardware and software requirements 479
 - highlights 476
 - key features of models 1RU and 1RX 477
 - key features of models 2RU and 2RX 478
 - physical characteristics 479
- FAStT600 Storage Server 473
 - description 474
 - devices supported 475
 - hardware and software requirements 475
 - highlights 473
 - limitations 475

- physical characteristics 475
- FAStT700 Storage Server 470
 - description 471
 - hardware requirements 472
 - highlights 470
 - limitations 472
 - physical characteristics 473
 - rack requirements 472
 - software requirements 472
- FAStT900 Storage Server 465
 - description 466
 - devices supported 469
 - hardware and software requirements 469
 - highlights 466
 - limitations 469
 - physical characteristics 469
- feature codes
 - 0920 AIX 5L Version 5.1 Bonus pack 589
 - 0921 AIX 5L Version 5.1 Expansion pack 589
 - 0948 AIX 5L Version 5.2 Bonus pack 589
 - 0949 AIX 5L Version 5.2 Expansion pack 589
 - 2623 4.7 GB SCSI-2 DVD-RAM Drive 434
 - 2624 32X (Max) SCSI-2 CD-ROM Drive 434
 - 2628 48X (Max) SCSI-2 Internal Auto-docking CD-ROM Drive 435
 - 2629 4.7 GB SCSI-2 Auto-docking DVD-RAM Drive 436
 - 2633 IDE CD-ROM Drive (Black bezel) 437
 - 2634 16X/48X(max) IDE DVD-ROM Drive 437
 - 2635 16X/48X(max) SCSI Auto-docking DVD-ROM Drive 438
 - 2640 IDE Slimline DVD-ROM Drive 439
 - 2826 GXT4000P Graphics Accelerator 419
 - 2827 GXT6000P Graphics Accelerator 423
 - 2842 GXT4500P Graphics Accelerator 421
 - 2843 GXT6500P Graphics Accelerator 425
 - 2848 GXT135P Graphics Accelerator 417
 - 2947 ARTIC960Hx 4-port Selectable Adapter 393
 - 2969 Gigabit Ethernet-SX PCI Adapter 400
 - 4023 SP Switch MX2 Adapter 361
 - 4025 SP Switch2 Adapter 361
 - 4026 SP Switch2 MX2 Adapter 361
 - 4032 SP Switch2 Interposer 362
 - 4217 6091 Cable 6ft 15-D-to-5BNC 428
 - 4237 Cable P201 15-D-to-13W3 - DDC/1010 428
 - 4238 CABLE, 15P D-SHELL TO 13W3, W/DDC, FOR P70/200 428

- 4239 15-D ADAPTER TO 3BNC ID CABLE 428
- 6120 IBM 80/160 GB Internal Tape Drive with VX 440
- 6131 60/150 GB 16-bit 8mm Internal Auto-docking Tape 440
- 6134 60/150 GB 16-bit 8mm Internal Tape Drive 441
- 6158 20/40GB 4mm Internal Tape Drive 441
- 6169 IBM 80/160 GB Internal Auto-docking Tape Drive 442
- 6185 20/40 GB 4mm Internal Auto-docking Tape Drive 443
- 6206 PCI single-ended Ultra SCSI Adapter 918
- 6206 PCI single-ended Ultra SCSI adapter 919, 924
- 6310 ARTIC960RxD Quad Digital Trunk Adapter 394
- 8398 SP Switch2 PCI-X Attachment Adapter 350
- Features of the 7040 pSeries 670/690 295
- FFDC, First-Failure Data Capture 14
- FFDC,7040 225
- Fibre Channel adapters 512
- Firmware 971
- firmware download 971
- First Failure Data Capture,7040 225
- FlashCopy Version 1,Enterprise storage server 2105 Model 800 463
- FlashCopy Version 2,FlashCopy Version 1,Enterprise storage server 2105 Model 800 463
- Flat Panel Color Monitor 325
- flat panel console 325
- flexible and scalable storage,2104 459
- floating-point execution unit,7040 224
- floating-point operation,7040 224
- FLOP,7040 224
- FMA, floating-point multiply/add 4
- four terms regarding memory 697
- Free Software Foundation (FSF) 592
- front door and front trim kit 319
- FTP 592
- Full System Partition 696

- G**
- General Parallel File System (GPFS) 674
- General Public License (GPL) 590
- generic data management tunnel 574
- Generic Security Service Application Program inter-
faces (GSS-API) 576
- Geodesic Systems Great Circle Versison 6.0.1.5 581
- Geographic Cluster Manager (GeoManager) 671
- Geographic Messaging (GeoMessage) 671
- Geographic Mirroring (GeoMirror) 671
- GeoRM 663
 - functional components 671
 - hardware requirement 666
 - software requirement 667
- GeoRM Version 2.4 670
- GNU Project 590, 592
- Gopher 592
- granularity 687
- graphical user interface on the HMC,7040 241
- graphics accelerators
 - available graphics accelerators 417
 - display and cable matrix 427
 - hardware classes 416
 - SoftGraphics product 416
 - withdrawn graphics accelerators 429
- GrAPHIGS, graphics accelerator API support 420
- GX bus slot,7040 220
- GX controller 9
- GXT120P, graphics accelerator 523
- GXT130P, graphics accelerator 523
- GXT135P, graphics accelerator 417
 - features 417
 - hardware requirement 418
 - software requirement 418
- GXT2000P, graphics accelerator 523
- GXT3000P, graphics accelerator 523
- GXT300P, graphics accelerator 523
- GXT4000P, graphics accelerator 419, 523
 - features 419
 - hardware requirement 420
 - software requirement 421
- GXT4500P, graphics accelerator 421
 - features 422
 - hardware requirement 423
 - software requirement 423
- GXT6000P, graphics accelerator 423, 523
 - features 424
 - hardware requirements 425
 - software requirement 425
- GXT6500P, graphics accelerator 425
 - features 426
 - hardware requirement 427
 - software requirement 427

H

HACMP

- basic concepts 628
- cluster configurations 634
 - concurrent disk access configurations 638
 - hot-standby configuration 634
 - mutual takeover configuration 635
 - non-concurrent disk access configurations 634
 - rotating standby configuration 635
 - third-party takeover configuration 636
- components 629
 - boot adapters and IP address takeover 630
 - clients 631
 - network adapters 630
 - networks 629
 - nodes 629
 - persistent node address 631
 - private network 629
 - public network 629
 - serial network 629
 - service network adapter 630
 - shared external disk devices 629
 - standby network adapter 630
- ESS/PPRC mirroring 663
- features 638
- features and benefits 639
- hardware supports matrix 652
- IP address takeover (IPAT) 630
- PSSP 662
- resource groups 632
 - cascading 633
 - concurrent 633
 - rotating 633
- router support 660
- RSCT 661–662
- SDD, Subsystem Device Driver 663
- software requirements 660
- supported communication adapters 655
- supported external storage subsystems 658
- supported rack-mounted storage subsystems 660
- supported systems 652
- HACMP and HACMP ES Version 4.2.1 644
- HACMP and HACMP ES Version 4.2.2 645
- HACMP and HACMP ES Version 4.3 646
- HACMP and HACMP ES Version 4.3.1 646
- HACMP and HACMP ES Version 4.4 647
- HACMP and HACMP ES Version 4.4.1 648
- ATM Classic IP and ATM LANE 648
- ATM Network Adapter 648
- OEM Disk API 648
 - software requirements 660
- HACMP and HACMP ES Version 4.5 649
- HAView and HATivoli 649
- IBM Generic Parallel File System (GPFS) 650
 - software requirements 661
- HACMP classic 639
 - Concurrent Resource Manager (CRM) 639
 - High Availability Network File System (HANFS) 639
 - High Availability Subsystem (HAS) 639
- HACMP Enhanced Scalability (ES) 639
- HACMP History 640
- HACMP Version 1.1 641
- HACMP Version 1.2 641
- HACMP Version 2.1 642
- HACMP Version 3.1 642
- HACMP Version 3.1.1 643
- HACMP Version 4.1 643
- HACMP Version 4.1.1 643
- HACMP Version 4.2.0 643
- HACMP Version 5.1 650
 - enhanced concurrent mode 651
 - General Parallel Files System (GPFS) 651
 - HACMP/XD (Extended Distance) 651
 - software requirements 662
- HACMP/XD (Extended Distance) 651
- HAGEO 663
 - configuration examples 667
 - functional components 671
 - Geographic Cluster Manager (GeoManager) 671
 - Geographic Messaging (GeoMessage) 671
 - Geographic Mirroring (GeoMirror) 671
 - asynchronous 671
 - synchronous 671
 - synchronous with Mirror Write Consistency 671
 - hardware requirement 665
 - releases and features 664
 - software requirement 667
- HAGEO Version 2.4 670
- half I/O drawer 233
- half I/O drawer,7040 218
- Hardware Management Console (HMC)
 - configuration notes 734
- Hardware Management Console(HMC)

- description 730
- Model highlights 7315-C02 730
- operating environment 733
- physical Specifications 733
- power requirements and thermal output 734
- technical overview 733
- hierarchical storage management (HSM) 537
- High Availability Network File System (HANFS) 639
- High Availability Subsystem (HAS) 639
- high performance computing,7040 227
- High-end pSeries 215
- High-end pSeries 670/690 215
- High-Performance adapters, adapter placement guidelines 904
- HMC 727, 729
 - user interface 728
- HMC graphical user interface 728
- HMC recovery software Version 1.3 562
- Host Bus Adapters (HBA) 568
- hot spare disk support 538
- hot spot management 538
- HPC feature 223
- HPC,7040 227
- hypervisor 600, 701, 716
- hypervisor call 701

I

- I/O book,7040 220
- I/O drawer
 - Model 6M2 190
- I/O drawer 7311-D10
 - features 209
 - hardware minimum requirements 208
 - I/O drawer attachment 208
 - maximum number 207
 - PCI slots 207
- I/O drawer 7311-D20 200
 - configuration notes 204
 - disk drive bays 203
 - features 209
 - hardware minimum requirements 204
 - internal disk drives 203
 - maximum number 204
 - PCI slots 203
 - remote I/O cables 204–205
 - RIO ports 204
 - RIO-2 ports 204
 - SPCN cables 204, 206
- I/O drawer attachment
 - Model 6C4 204
 - Model 6M2 205
- I/O drawer,7040-61D 232
- I/O planar,7040 233
- IBF,7040 239
- IBM AIX Developer Kit, JAVA2 Technology 542, 570, 580
- IBM Directory 562
- IBM Directory Server 574, 579
- IBM external disk storage 452
- IBM Global Security Toolkit Version 4.0 577
- IBM Global Security Toolkit Version 5.0 577
- IBM Grid Toolbox for AIX 576
- IBM Hardware Management Console for pSeries 727, 729
- IBM Hardware Management Console for pSeries,7040 240
- IBM HTTP Server 562, 579
- IBM IP Security Version 5.1 and Version 5.2 578
- IBM Just-In-Time (JIT) compiler 542
- IBM Linux support line 591
- IBM Linux Technology Center (LTC) 591
- IBM Network Authentication Service Version 1.3 576
- IBM Self-Voicing Kit (SVK) 529
- IBM Solution Partnership Centers 591
- IBM Tivoli Storage Manager Version 5.1.5 583
- IBM Tivoli Storage Resource Manager Version 1.2 583
- IBM Total Storage 451
- ICU, Internationalized Classes for Unicode 532
- IDE 437
- Interaction of AIX and firmware in a partition 709
- Internal battery feature,7040 239
- internal battery feature,7040 239
- internal disk bays
 - Model 6M2 176
- internal disk drives
 - Model 6M2 176
- internal tape drives 440
- interoperability 582
- Inventory Scout 545
- IP fragmentation 549
- IP multi-routing and multiple gateway 546
- ISO 13406-2, graphics accelerator 422
- ISO 9241, graphics accelerator 420
- ISO Standard ISO8859-15 532

J

Java 2 Enterprise Edition (J2EE) Version 1.3 579
Java Authentication and Authorization Service (JAAS) Version 1.0 581
Java Cryptography Extension (JCE) 543, 580
Java Messaging Server (JMS) 579
Java Secure Sockets Extension (JSSE) 543, 581
Java security 542
Joint Test Action Group,7040 225
Journaled File System (JFS) 2 538
JTAG,7040 225

K

Kerberos Version 5 552
Kernel Debugger (KDB) 541
key distribution center (KDC) 576

L

L3 cache
 Model 6M2 176
L1 cache
 Model 6M2 176
L2 cache 345
 Model 6M2 176
L2 cache subsystem,7040 224
L2 cache,7040 224
L3 cache 345
L3 cache controller and directory,7040 225
L3 controller,7040 225
L3 directory,7040 225
Large page support 710
Lawrence Berkeley National Laboratory (LBNL) 548
LDAP, Lightweight Directory Access Protocol 525
LED codes 732
lg_dumplv 541
lighting and transforms, graphics accelerator 416
Linus Torvalds 589
Linux 589
 AIX Toolbox for Linux application 604
 configuration examples 598
 Express Configurations 602
 fundamental components 592
 General Public License (GPL) 590
 GNU Project 590, 592
 history of Linux 589
 I/O Device and adapter support 597
 IBM Linux support line 591

IBM Linux Technology Center (LTC) 591
Kernel 593
Linus Torvalds 589
Linux affinity on AIX 5L 602
Logical partitioning (LPAR) 600
Open Source Development Lab 591
Open Source Software in the AIX Toolbox 605
other sources of Open Source Software 606
performance considerations 603
pSeries RAS 599
pSeries scalability 599
Red Hat Linux 594
software service and support 607
supported pSeries systems 595
SuSE Linux 594
TurboLinux 595
United Linux 594
Linux affinity 545, 571, 602
LMB 700
Loadable Authentication Module Framework (LAMF) 573
Logical 695
logical memory 699
Logical Memory Block 700
logical partition 687
logical partitioned environment 695
Logical partitioning (LPAR) 600
Logical partitioning overview 686
Logical partitioning support on pSeries 630 and pSeries 650 695
logical track group (LTG) 537
Lotus Domino Enterprise 5.0.2b Webserver 562
LPAR configuration notes 690
 pseries 630 Model 6C4/6E4 693
 pseries 650 Model 6M2 692
 pseries 655 Model 651 692
 pseries 670 Model 671 691
 pseries 690 Model 681 691
lsattr 707
LTO 491
 models 491
 product positioning 491

M

Managing partitions 727
Maximum number of processors, memory size, and partitions 689
MBCS, Multi-Byte Character Set Support 529

- MCAD 421
- MCAE 421
- MCM population order 229
- MCM, Multi-chip module 344
- MDA,7040 240
- Mechanical Computer Aided Design 421
- media bays
 - Model 6M2 176
- Media drawer 236
- memory
 - Model 6M2 187
- memory subsystem for pSeries 670 228
- microprocessors
 - 32-bit 3
 - 64-bit 3
 - copper chip wiring 2
 - P2SC 5
 - POWER architecture 4
 - POWER1 4
 - POWER3 6
 - POWER4 1, 7
 - POWER4+ 1, 10
 - PowerPC 5
 - processor descriptions 3
 - RS64 6
 - RS64-II 6
 - RS64-III 6
 - RS64-IV 6
 - silicon-on-insulator (SOI) 8
- Mirror Write Consistency (MWC) 537
- MLD. Merged Logic DRAM 7
- Mobile IPv6 572
- Model 140
 - site and hardware planning information 749
- Model 150
 - site and hardware planning information 749
- Model 150 multiple graphics adapter placement guide 864
- Model 150,adapter placement guidelines 859
- Model 170
 - site and hardware planning information 775
- Model 170 multiple graphics adapter placement guide 874
- Model 170,adapter placement guidelines 867
- Model 265,adapter placement guidelines 897
- Model 270 multiple graphics adapter placement guide 882
- Model 270,adapter placement guidelines 876
- Model 651,adapter placement guidelines 850
- Model 671 242
 - configuration notes 245
 - Capacity Upgrade on Demand for memory 264
 - Capacity Upgrade on Demand for processors 262
 - Central Electronics Complex (CEC) 245
 - clustered server environment 260
 - disks and boot devices 254
 - dynamic logical partitioning 265
 - Hardware Management Console for pSeries (HMC) 249
 - I/O drawer attachment 252
 - I/O drawers 251
 - Logical partitioning (LPAR) 250
 - memory 247
 - On/Off Capacity on Demand for processors 263
 - PCI and PCI-X slots and adapters 256
 - physical specifications/power and operating requirement 266
 - rack and power 255
 - Trial Capacity on Demand (Trial CoD) 264
 - highlights 243
 - minimum and standard features 244
 - system expansion 244
- Model 671,adapter placement guidelines 824
- Model 681 266
 - configuration notes 269
 - Capacity Upgrade on Demand for memory 294
 - Capacity Upgrade on Demand for processors 292
 - Central Electronics Complex (CEC) 269
 - clustered server environment 290
 - disks and boot devices 283
 - dynamic logical partitioning 295
 - Hardware Management Console for pSeries (HMC) 277
 - I/O drawer attachment 281
 - I/O drawers 279
 - logical partitioning (LPAR) 279
 - memory 271
 - On/Off Capacity on Demand for processors 293
 - PCI and PCI-X slots and adapters 286
 - physical specifications/power and operating requirement 295
 - rack and power 283

- Trial Capacity on Demand (Trial CoD) 293
 - highlights 266
 - minimum and standard features 267
 - system expansion 268
- Model 681, adapter placement guidelines 835
- Model 6C1, adapter placement guidelines 803
- Model 6C3, adapter placement guidelines 820
- Model 6C4
 - 7311-D20
 - maximum number 204
 - remote I/O cables 204
 - SPCN cables 204
 - I/O drawer 200
 - I/O drawer 7311-D20
 - maximum number 204
 - I/O drawer attachment 204
- Model 6C4 (4-Slot PCI Riser), adapter placement guidelines 809
- Model 6C4 (6-Slot PCI Riser), adapter placement guidelines 815
- Model 6E1, adapter placement guidelines 803
- Model 6E3, adapter placement guidelines 820
- Model 6E4 (4-Slot PCI Riser), adapter placement guidelines 809
- Model 6E4 (6-Slot PCI Riser), adapter placement guidelines 815
- Model 6F0, adapter placement guidelines 789
- Model 6F1, adapter placement guidelines 789
- Model 6M2
 - 7311-D10
 - hardware minimum requirements 208
 - I/O drawer attachment 208
 - maximum number 207
 - PCI slots 207
 - 7311-D20
 - disk drive bays 203
 - hardware minimum requirements 204
 - internal disk drives 203
 - maximum number 204
 - PCI slots 203
 - SPCN cables 206
 - AIX operating system 177
 - boot devices 191
 - Capacity Upgrade on Demand 177
 - CD-ROM 176
 - configuration notes
 - recommended memory configuration 188
 - disk 191
 - diskette drive 177
 - DVD-RAM 176
 - DVD-ROM 176
 - Express Configurations 195
 - features 179
 - highlights 173
 - hot-plug 194
 - I/O drawer 190, 200
 - I/O drawer 7311-D10
 - hardware minimum requirements 208
 - highlights 207
 - I/O drawer 7311-D20
 - disk drive bays 203
 - hardware minimum requirements 204
 - internal disk drives 203
 - maximum number 204
 - PCI slots 203
 - remote I/O cables 205
 - SPCN cables 206
 - I/O drawer attachment 190, 205
 - integrated Ethernet adapter 177
 - integrated SCSI adapters 177
 - internal disk bays 176
 - internal disk drives 176
 - L1 cache 176
 - L2 cache 176
 - L3 cache 176
 - logical partitioning 194
 - media 191
 - media bays 176
 - memory 187
 - minimum features 176
 - On/Off Capacity on Demand 177
 - PCI slots 192
 - positioning 172
 - power 189
 - processors 187
 - racks 189
 - RAS features 175
 - standard ports 177
 - system architecture 174
 - system expansion 178
 - disk drive 178
 - memory 178
 - PCI expansion slots 178
 - technical overview 173
 - GX bus 175
 - RIO-2 173
 - Single Chip Module 175
 - split-backplane 174

- warranty 177
 - Model 6M2 I/O drawer 7311-D10
 - I/O drawer attachment 208
 - maximum number 207
 - PCI slots 207
 - Model 6M2,adapter placement guidelines 845
 - Model 7014-T00 310
 - AC power distribution bus 313
 - description 311
 - front door and front trim kit 312
 - hardware requirements 316
 - limitations 314
 - model highlights 7014-T00 311
 - other standard equipment and optional features 312
 - physical specifications, operating environment and power requirement 312
 - publications 317
 - ruggedized rack feature 312
 - side panels and side to side rack suite attachment 311
 - Type 6 power distribution bus 313
 - Type 7 power distribution bus 313
 - Model 7014-T42 317
 - AC power distribution bus 320
 - description 318
 - front door and front trim kit 319
 - hardware requirements 324
 - limitations 323
 - Model highlights 7014-T42 318
 - other standard equipment and optional features 320
 - physical specifications, operating environment and power requirement 320
 - ruggedized rack feature 319
 - side panels and side to side rack suite attachment 319
 - Type 6 power distribution bus 320
 - Type 7 power distribution bus 321
 - Model 7316-TF2 325
 - physical specification 326
 - Model 9910-A08 331
 - Model 9910-A10 331
 - Model 9910-A13 332
 - Model 9910-A30 333
 - Model 9910-Axx 330
 - Model 9910-M10 339
 - Model 9910-M11 339
 - Model 9910-M15 340
 - Model 9910-M30 341
 - Model 9910-Mxx 338
 - Model 9910-P07 335
 - Model 9910-P18 335
 - Model 9910-P22 336
 - Model 9910-P33 337
 - Model 9910-P60 338
 - Model 9910-Pxx 333
 - Model B50,adapter placement guidelines 884
 - Model B80,adapter placement guidelines 798
 - Model D10,adapter placement guidelines 887
 - Model D20,adapter placement guidelines 891
 - Model F40
 - facts and features summary 21–23
 - Model F50
 - facts and features summary 21–23
 - Model H70
 - site and hardware planning information 756, 762, 768
 - Model M2
 - 7311-D20
 - remote I/O cables
 - I 205
 - Models 6F0 and 6F1,adapter placement guidelines 789
 - Modular I/O (MIO) Library Version 2.3 581
 - Motor Drive Assembly,7040 240
 - Motor/Scroll Assembly,7040 240
 - MRPD, Machine Reported Product Data 545
 - MSA,7040 240
 - MSR.DR 702
 - MSR.IR 702
 - multichip module,7040 222, 226
 - Multimedia 582
 - Multi-node Load Balanced (MNLB) 546
 - multiprotocol over ATM (MPOA) 549
- ## N
- Netscape Communicator 4.8 580
 - Netscape Enterprise Server 3.6.3, 4.0 562
 - Netscape FastTrack Server 3.01 562
 - Network Buffer Cache (NBC) 546
 - network installation management (NIM) 564
 - network interface backup mode 547
 - networking security 574
 - new firmware download 971
 - NFS statd multithreading 533
 - noise emission notes 783

NotesBench 953
nPartitions 687

O

On demand operating environment 13
On/Off Capacity on Demand (CoD) 16
Open Secure Shell (OpenSSH) Version 3.6 578
Open Source Development Lab 591
OpenGL, graphics accelerator API support 420
operating systems
 AIX 5L Version 5.1 522
 AIX 5L Version 5.2 555
 Linux 589

P

P2SC microprocessor 5
p655 344
packet capture library 548
page frame table 698
pages 697
Paging performance in a partitioned environment 709
paging performance in a partitioned environment 709
Parallel ESSL 673
Parallel System Support Program (PSSP) 673
partition and system profiles 712
Partition isolation and security 696
Partition on Demand 711
partition profile 712
partition reset capabilities 707
partition resource
 I/O slots 714
 memory 714
 processors 714
partition resources 712
partitioned environment
 Fast reboot 709
 Paging performance 709
Partitioning 686
Partitioning implementation
 interrupt controller 702
 PCI host bridges 702
partitioning implementation 701
 boot devices 719
 Capacity Upgrade on Demand 711
 console device 723
 CUoD 711

EEH 724
error handling 703
firmware 703
graphics console 723
hardware 701
hardware password 705
high availability 724
HPC 710
hypervisor 703
interrupt controller 702
lscfg 705
media devices 718
native ISA devices 722
network devices 722
operating system
 AIX 5L Version 5.1 707
 AIX 5L Version 5.2 710
PAP 705
PCI host bridge 702
PCI host bridges 702
PHB 724
Recommended Maintenance Level
 5100-01 708
recommended Maintenance Level
 5100-02 710
RTAS 705
serial console 723
service processor 703
SMS 705
systems management services 705
Time-of-day 707
TOD 707
Partitioning implementation on pSeries servers 701
 Media devices 718
 Network devices 722
partitioning implementation on pSeries servers
 I/O device assignment considerations 717
Partitioning support on pSeries servers 688
 supported models 688
partitioning-capable pSeries server 688
PartnerWorld for developers 582
PASSTHRU Support 529
PDU Combinations for 7014 Models T00 and T42 316
Peer-to-Peer Remote Copy (PPRC) 651
Peer-to-Peer Remote Copy (PPRC) Version 1, Enterprise storage server 2105 Model 800 464
Peer-to-Peer Remote Copy (PPRC) Version 2, Enterprise storage server 2105 Model 800 464

- Performance 950
- Performance Benchmarks 950
- Perl language 571
- per-partition NVRAM 707
- Per-partition Time-Of-Day clock values 707
- persistent node address 631
- PFT 698
- PHB 702
- physical memory 698
- Physical Memory Block 700
- physical partition 687
- pipeline structure 224
- PKCS, Public-Key Cryptography Standards 554
- Pluggable Authentication Mechanism (PAM) 553
- Pluggable Authentication Modules (PAM) 573
- PMB 700
- POR 225
- Portable Document Format (PDF) 582
- positioning pSeries 670/690 216
- Post-10/2002 firmware 716
- power
 - Model 6M2 189
- POWER 4 microprocessor
 - copper interconnects 7
- POWER architecture 4
- Power cord features 907
 - features codes table 907
 - power cord connector spotters reference 911
- Power On Reset,7040 225
- power subsystem,7040 238
- POWER1 microprocessor 4
- POWER2 227
- POWER3 microprocessor 6
- POWER4 chip,7040 222
- POWER4 core 224
- POWER4 microprocessor 1, 7
 - CIU, Core Interface Unit 8
 - First Failure Data Capture 9
 - GX controller 9
 - JTAG interface 10
 - pervasive function 9
 - Power On Reset (POR) 9
 - silicon-on-insulator (SOI) 2
- POWER4 processor and MCM packaging 222
- POWER4+ microprocessor 1, 10
- PowerPC microprocessor 5
- PPRC Version 1 464
- PPRC Version 2 464
- Pre-10/2002 firmware 716
- primary I/O book,7040 230
- private network 629
- processor bus pass through module,7040 229
- processors
 - Model 6M2 187
- Protection against inter-partition data access 696
- pSeries 610 Models 6C1 and 6E1,adapter placement guidelines 803
- pSeries 615 Models 6C3 and 6E3,adapter placement guidelines 820
- pSeries 630 Models 6C4 and 6E4 (4-Slot PCI Riser), adapter placement guidelines 809
- pSeries 630 Models 6C4 and 6E4 (6-Slot PCI Riser), adapter placement guidelines 815
- pSeries 640 Model B80,adapter placement guidelines 798
- pSeries 650 Model 6M2,adapter placement guidelines 845
- pSeries 655 344
- pSeries 655 Model 651,adapter placement guidelines 850
- pSeries 670 215, 242
- pSeries 670 media drawer 236
- pSeries 670 Model 671 drawers 61D,adapter placement guidelines 824
- pSeries 670,I/O books 230
- pSeries 670/690 215
- pSeries 690 215, 266
- pSeries 690 media drawer 236
- pSeries 690 Model 681 drawer 61D,adapter placement guidelines 835
- pSeries 690,I/O books 230
- pSeries fibre channel Host Bus Adapters (HBA) 512
 - feature 6228 512
 - feature 6239 512
- pSeries firmware download 971
- pSeries Performance
 - NotesBench 953
 - rPerf, Relative Performance 952
 - SPEC, System Performance Evaluation Corporation 952
 - TPC, Transaction Processing Council 952
- Public Key Cryptography Standards (PKCS) 543, 581
- Public Key Infrastructure (PKI) 573
- public network 629

Q

QBB 687
Quad Building Block 687
Quality of Service (QoS) 549, 573

R

Rack Tray for monitor 325
rack-mounted flat panel console 325
racks
 6M2 189
Racks and rack solution 309
 Model 7014-T00 overview 310
 Model 7014-T42 overview 317
 Model 7316-TF2 overview 325
RAN 391
RAN breakout box 731
RAS features
 Model 6M2 175
RAS, Reliability, Availability, and Serviceability 539, 569
RAS, Reliability, availability, and serviceability 14
RAS,7040 218
rasterization, graphics accelerator 416
RDBMS 710
real memory 699
real mode limit register 702
real mode offset register 702
Red Hat Linux 594
Redbooks Web site 991
 Contact us xxxvii
register rename pool,7040 224
Remote Asynchronous Node 391
reserved memory regions in a partitioned environment 715
Resource
 addition, DLPAR 725
 movement, DLPAR 725
 removal, DLPAR 725
resource groups 632
resource management 685
resource value
 Desired 715, 726
 Maximum 715
 Minimum 715
 Required 715, 726
Richard Stallman 593
RIO connector,7040 233
RIO drawer,7040 234–235

RIO Riser card,7040 233
RIO-2 351
RIO-2 connector,7040 233
RIO-2 drawer,7040 234, 236
RML 702
RMO 702
rotating resource groups 633
rPerf, Relative Performance 952
RS/6000 firmware download 971
RS/6000 SP 359
RS-422 cable 731
RS64 microprocessor 6
RS64-II microprocessor 6
RS64-III microprocessor 6
RS64-IV microprocessor 6
RSCT, Reliable Scalable Cluster Technology 526, 577
RSCT, Reliable, scalable cluster technology 12
Runtime Predictive Deconfiguration of Processors (RPDP) 566

S

SAN Data Gateways 513
SAN solution 511
 IBM TotalStorage SAN switches 513
 Large configuration 519
 Medium configuration 518
 Sample configurations 517
 Small configuration 517
SAN switches 513
 2109-F16 515
 2109-F32 515
 2109-M12 516
 3534-F08 513
SBCS, Single Byte Character Set Support 529
SCSI cabling examples 917
 PCI differential-ended Ultra SCSI adapter cables and terminators
 device-to-device cables 926
 terminators 927
 cabling examples 928
 PCI Single-Ended Ultra SCSI adapter 918
 PCI single-ended Ultra SCSI adapter cables and terminators 919
 device-to-device cables 920
 Single-Ended adapter-to-first device cables 920
 system-to-system cables 921

- terminators 922
- cabling examples 922
- supported Ultra configurations 919
- SDD, Subsystem Device Driver 663
- secondary I/O books,7040 230
- Secure Multi-Purpose Internet Mail Extensions (S/MIME) 581
- Secure Sockets Layer (SSL) 543
- SecureWay Directory Server 524
- segments 697
- Sendmail Version 8.11.0 548
- serial network 629
- server load balancing (SLB) 546
- Service Agent 735
- Service authority 724
- service authority 724
- service network adapter 630
- service processor 732
- serviceability 218
- several partitioning implementations 686
- Side panels and side to side rack suite attachment 319
- Simple Network Management Protocol (SNMP) 571
- single core POWER4 processor feature 227
- site and hardware planning guide 747
 - 7014 Model T00 and T42 782
 - 7026 Model B80 748
 - 7028 Model 6C1 and 6E1 750
 - 7028 Model 6C4 and 6E4 751
 - 7029 Model 6C3 and 6E3 753
 - 7038 Model 6M2 754
 - 7039 pSeries 655 756
 - 7040 pSeries 670 762
 - 7040 pSeries 690 768
 - 7043 43P Model 150 773
 - 7044 44P Model 170 775
 - 7311 Model D10 777
 - 9112 Model 265 779
 - 9114 Model 275 780
- site and hardware planning Information 747
- SMP, symmetric multiprocessor 5
- SNMPv3 Encryption Version 5.2 578
- SOI CMOS technology 224
- SOI, silicon-on-insulator 2, 8
- SP Switch2 349
- SPEC, System Performance Evaluation Corporation 951
- standby network adapter 630

- Storage Networking Industry Association (SNIA), 568
- storage resource management (SRM) 583
- Storage, communication and other I/O adapters 365
 - ARTIC adapters 393
 - asynchronous adapters 390
 - asynchronous transfer mode adapters 398
 - available PCI I/O adapters 365
 - cryptographic adapters 411
 - Ethernet and token-ring adapters 400
 - PCI Storage Adapters 368
 - SCSI and SSA Overview
 - general terminology 448
 - SCSI 444
 - SCSI-I 444
 - SCSI-II 444
 - SCSI-III 445
 - SCSI-III Standards 445
 - system adapters 386
 - WAN adapters 397
- superscalar 224
- Supervisor 702
- Supported partitioning-capable pSeries servers 688
- Supported Peripherals by Device Matrix 931
 - Disk drives and subsystems 932, 939
 - Displays 934, 941
 - Expansion Cabinets 938, 945
 - Graphics related 934, 941
 - Optical drives and libraries 935, 942
 - Storage Network 935, 942
 - Switch 932, 938
 - Tape drives and libraries 936, 943
- SuSE Linux 594
- System life time information
 - 7038-6M2 IBM pSeries 650 Server Model 6M2 964
 - 7039-651 IBM pSeries 655 Server Model 651 965
 - 7046-B50 IBM RS/6000 Server Model B50 966
 - 7202-900 IBM RS/6000 Expansion Rack Model 900 966
 - 7236-001 IBM RS/6000 Media Streamer Model 001 967
 - 7315-C01 IBM Hardware Management Console C01 968
 - overview 7006 IBM RS/6000 956
 - overview 7007 IBM RS/6000 956

- overview 7008 IBM RS/6000 956
- overview 7009 IBM RS/6000 957
- overview 7010 IBM RS/6000 957
- overview 7011 IBM RS/6000 957
- overview 7012 IBM RS/6000 958
- overview 7013 IBM RS/6000 959
- overview 7014 IBM RS/6000 960
- overview 7015 IBM RS/6000 960
- overview 7017 961
- overview 7024 IBM RS/6000 Server 962
- overview 7025 IBM RS/6000 Server 962
- overview 7026 IBM RS/6000 Server 963
- overview 7027 IBM High Capacity Drawer 963
- overview 7028 IBM pSeries 610 Server 964
- overview 7030 IBM RS/6000 964
- overview 7040 IBM pSeries 670/690 965
- overview 7043 IBM RS/6000 43P 965
- overview 7044 IBM RS/6000 966
- overview 7248 IBM RS/6000 43P Series 967
- overview 7311 IBM Rack-Mounted Drawers 967
- overview 7316 IBM Rack-Mounted Console 968
- overview 7317 IBM RS/6000 Telecom Server and External SCSI Disk 968
- overview 7318 IBM Serial Communications Network Server 969
- overview 7319 IBM RS/6000 Fibre Channel Switch and Adapter 969
- System profile 712
- System rack 7040-61R 237
 - cooling 239
 - internal battery feature 239
 - power subsystem 238

T

- tape drive products 479
- tape libraries 485
- TCE 702
- TCP splicing 547
- texture mapping, graphics accelerator 416
- Three assignable resource types 714
- three kinds of values for resource assignment 715
- Ticket Granting Ticket (TGT) 553
- TLB 698
- Token Ring support for MPOA 549
- TotalStorage Enterprise Storage Server (ESS) 461
- TotalStorage FAStT200 Storage Server 476
- TotalStorage FAStT600 Storage Server 473
- TotalStorage FAStT700 Storage Serve 470
- TotalStorage FAStT900 Storage Server 465
- TPC, Transaction Processing Council 952
- translate-off 699
- translation control entry 702
- translation look aside buffer 698
- translation-off mode 702
- Transport Layer Security (TLS) protocols 543
- Trial Capacity on Demand 17
- TurboLinux 595
- Type 1722-60U 473
- Type 1742 -90U 465
- Type 1742-1RU 470
- Type 3542-1RU 476
- Type 3542-1RX 476
- Type 3542-2RU 476
- Type 3542-2RX 476
- Type 6 power distribution bus 313
- Type 7 power distribution bus 313

U

- Unexpected partition crash 696
- Unicode support 530
- Uninterruptible Power System (UPS) 326
 - 9910-A08 331
 - 9910-A10 331
 - 9910-A13 332
 - 9910-A30 333
 - 9910-M10/M11 339
 - 9910-M15 340
 - 9910-M30 341
 - 9910-P07 335
 - 9910-P18 335
 - 9910-P22 336
 - 9910-P33 337
 - 9910-P60 338
- battery runtime 329
- choosing the right MGE UPS 339
- choosing the right Powerware UPS 334
- devices supported 329
- rack rules 328
- UPS Model 9910-Axx 330
- UPS Model 9910-Mxx 338
- UPS Model 9910-Pxx 333
- Uninterruptible Power System battery runtime 329
- United Linux 594
- Universal Disk Format (UDF) 567

unix_64 kernels 541
unix_mp 541
unix_up 541
UPS 326
UPS battery runtime 329

V

VERITAS Cluster Server 584
VERITAS File System 583
VERITAS Foundation Suite 583
VERITAS Volume Manager 583
Virtual IP Address (VIPA) 547
Virtual Local Area Network (VLAN) 549
virtual memory 697
virtual memory manager 698
Virtual Shared Disk (VSD) 673
virtual terminal device support 709
VMM 698
VXA 440

W

WAN adapters
 2943 8-port Asynchronous Adapter 390
 2944 128-port Asynchronous Adapter 391
Web servers 592
Web-based System Manager 722
Web-based System Manager Security Version 5.1
and Version 5.2 578
WebSphere Application Server 579
wide-area information server (WAIS) 592
Wireless Application Protocol (WAP) Gateway Ver-
sion 1.1 580
WLM, Workload Manager 535, 566
WSM application support 527
WSM, Web-based System Manager 526

X

X.509 standard 576
X11, graphics accelerator API support 420
XRC 464



Redbooks

IBM [®]SERVER pSeries Systems Handbook 2003 Edition

(1.5" spine)
1.5" x 1.998"
789 <-> 1051 pages



IBM [®]@server pSeries Systems Handbook 2003 Edition



The ideal deskside reference for the latest pSeries models and features

The *IBM @server pSeries Systems Handbook 2003 Edition*, SG24-5120 is a comprehensive single-source guide covering the pSeries product line. In some cases, withdrawn products are included for comparison. Major hardware and software offerings are introduced and their prominent functions discussed.

Hundreds of tables and figures to accelerate your research

This IBM Redbook is suitable for professionals wishing to acquire a better understanding of pSeries products, including customers, sales and marketing professionals, technical support professionals, and IBM Business Partners.

A required reading for all pSeries and AIX professionals

Inside this redbook, you will find a historical look at RS/6000 and pSeries hardware, an overview of the latest pSeries models, a short discussion on hardware architecture, information on storage, graphics, and communications features, a description of AIX and supported software platforms, and hundreds of tables and figures providing effective access to useful information.

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

This redbook does not replace the latest pSeries 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 solutions for the UNIX marketplace.

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