


RS/6000 7046 Model B50 Handbook

Scott Vetter, Tetsuroh Maeda



International Technical Support Organization

www.redbooks.ibm.com

SG24-7046-00



International Technical Support Organization

RS/6000 7046 Model B50 Handbook

December 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, "Special Notices" on page 181.

First Edition (December 1999)

This edition applies to the IBM RS/6000 7046 Model B50 ISP Server, AIX Version 4.3.3, program number 5765-C34, and Linux 2.2.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JN9B Building 003 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

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

© Copyright International Business Machines Corporation 1999. All rights reserved.

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

Contents

Figures	vii
Tables	ix
Preface	xi
The Team That Wrote This Redbook	xi
Comments Welcome	xii
Chapter 1. RS/6000 Introduction	1
1.1 Basic ISP Server Topology	2
1.2 RS/6000 7046 Model B50 Overview	3
1.2.1 RS/6000 7046 Model B50 Key Attributes	4
1.2.2 RS/6000 7046 Model B50 Feature Summary	5
1.3 Performance Results	6
1.3.1 SPEC95 Software Benchmark	6
1.3.2 Performance Summary	9
1.4 Industry Trends and Directions	10
Chapter 2. Detailed Product Description	11
2.1 IBM RS/6000 7046 Model B50 Product Description	11
2.1.1 Standard Features	12
2.1.2 Reliability, Availability, and Serviceability (RAS)	18
2.1.3 Supported Optional Features	20
2.1.4 RS/6000 7046 Model B50 Adapter Placement Guidelines	26
2.1.5 Service Package	26
2.1.6 Options Library (Hard Copy)	27
Chapter 3. Adapters: Overview and Configuration	29
3.1 Communication Adapters	29
3.1.1 Token-Ring Adapter	29
3.1.2 Ethernet Adapter	30
3.1.3 ISDN Adapter	30
3.1.4 FDDI Adapters	31
3.1.5 ATM Adapters	33
3.1.6 Multiple-Port Adapters	34
3.1.7 WAN Adapters	34
3.2 PCI Storage Adapters	35
3.2.1 SCSI Adapters	35
3.2.2 SSA Adapter	37
3.3 Graphics Accelerators	38
3.3.1 POWER GXT130P (# 2830)	38

3.4	Performance Notes for PCI Communications and Storage Adapters . . .	39
3.4.1	Adapter Placement	39
3.4.2	Networking Throughput	39
3.4.3	Disk I/O Throughput	44
3.5	Adapter and Device Configuration on PCI-Based RS/6000 Systems . . .	48
3.5.1	Device Types	48
3.5.2	Device Location Codes	49
3.5.3	PCI Adapter Configuration	54
3.6	Summary Chart of Adapters, Size, and Placement	56
Chapter 4. Storage: Architectures and Devices		59
4.1	Storage Boot Devices	59
4.2	SCSI Overview	59
4.2.1	SCSI-I	59
4.2.2	SCSI-II	60
4.2.3	SCSI-III	61
4.2.4	SCSI Terminology	64
4.2.5	SCSI Repeaters	66
4.2.6	Summary of SCSI Specifications	67
4.3	SCSI Cabling	67
4.3.1	General SCSI Cabling Considerations	68
4.3.2	Cabling the PCI Single-Ended Ultra SCSI Adapter	71
4.3.3	Cabling the PCI Differential-Ended Ultra SCSI Adapter	75
4.3.4	Cabling the PCI Dual-Channel Ultra2 SCSI Adapter (# 6205)	80
4.4	Serial Storage Architecture Overview	81
4.5	Comparison between SCSI and SSA Architecture	82
4.6	Internal SCSI Devices	83
4.6.1	Disk Drives for the Model B50	83
4.6.2	Internal CD-ROM Drive	84
4.7	Internal SSA Devices	85
4.8	External SCSI and SSA Storage Devices	85
4.8.1	Expandable Storage Plus (# 2104)	85
4.8.2	IBM 7133 Serial Disk System Advanced Models T40 and D40	87
Chapter 5. Hardware Architecture Overview		91
5.1	Hardware Design for the RS/6000 7046 Model B50	91
5.2	Bus Architecture	93
5.2.1	The PCI Bus Architecture	93
5.2.2	The ISA Bus Architecture	94
5.2.3	The System Buses	95
5.2.4	The 6XX-MX I/O Bus	95
5.2.5	The X-Bus	96
5.3	The Processor Subsystem	96

5.3.1	The PowerPC 604e Microprocessor	96
5.4	The RS/6000 Platform Architecture	98
5.4.1	The RS/6000 Platform Architecture Introduction	98
5.4.2	RS/6000 Platform Architecture Achievements	99
Chapter 6. Diagnostics and System Management Services		101
6.1	Diagnostics Operating Considerations	101
6.2	Standalone Diagnostics	103
6.2.1	Booting the Standalone Diagnostics CD-ROM	103
6.2.2	Running Standalone Diagnostics in Concurrent Mode	104
6.3	Online Diagnostics	104
6.3.1	Service Mode	105
6.3.2	Concurrent Mode	106
6.3.3	Single-User Mode	107
6.4	Diagnostic Tasks and Service Aids	108
6.5	Understanding the Firmware	112
6.5.1	General RS/6000 Boot Sequence	112
6.5.2	Typical Boot Sequence for RS/6000 7046 Model B50	113
6.6	Power-On Self-Test	114
6.6.1	Processor POST	114
6.6.2	Memory DIMM and L2 Cache POST	114
6.6.3	Graphics Adapter POST	116
6.6.4	Keyboard and Mouse Controller POST	116
6.6.5	SCSI Subsystem POST	116
6.7	System Management Services (SMS)	116
6.7.1	SMS Startup	117
6.7.2	SMS Graphical Main Menu	118
6.7.3	SMS ASCII Main Menu	123
6.7.4	SMS Firmware Update	124
6.7.5	SMS Firmware Checkpoints	127
Chapter 7. Application Software Solutions for the Model B50		129
7.1	Internet Services for ISP	130
7.1.1	Domain Name Service	131
7.1.2	Mail Service	131
7.1.3	Web Service	133
7.1.4	FTP Service	133
7.1.5	Directory Service	133
7.1.6	Security and System Management	133
7.2	Applications for Businesses	133
7.2.1	File Sharing Service	133
7.2.2	Printer Sharing Service	134
7.3	Applications for Internet Service	134

7.3.1 Basic Internet Features on AIX Version 4.3.3	134
7.3.2 Service Providers Utility Pack	136
7.3.3 IBM WebSphere Application Server Version 2.02	140
7.3.4 WebSphere Performance Pack Version 2.0	143
7.3.5 SecureWay Firewall for AIX V4	148
7.3.6 AIX Fast Connect for Windows	151
Chapter 8. Linux on RS/6000	153
8.1 What is Linux?	153
8.2 LinuxPPC for RS/6000	153
8.3 Support for LinuxPPC on RS/6000 Platform	154
8.4 Applications on LinuxPPC	155
8.4.1 Basic Features for Internetworking on Linux 2.2	155
8.4.2 The Apache HTTP Server	156
8.4.3 Squid - Firewall/Web Cache Server	156
8.4.4 SAMBA - SMB/CIFS Compatible File/Printer Server	157
8.5 Installing LinuxPPC on the Model B50	158
8.5.1 Installation Considerations	158
8.5.2 Installation Procedures	159
Appendix A. Special Notices	181
Appendix B. Related Publications	185
B.1 IBM Redbooks Publications	185
B.2 IBM Redbooks Collections	185
B.3 Other Resources	186
B.4 Referenced Web Sites	186
How to Get IBM Redbooks	189
IBM Redbooks Fax Order Form	190
Glossary	191
Index	197
IBM Redbooks Evaluation	205

Figures

1. RS/6000 7046 Model B50 Installed in S00 Rack	1
2. Typical ISP Topology	2
3. RS/6000 7046 Model B50	4
4. RS/6000 7046 Model B50	11
5. RS/6000 7046 Model B50 - System Board Locations	13
6. RS/6000 7046 Model B50 - Front View	15
7. RS/6000 7046 Model B50 - Rear View	16
8. Universal Mounting Rails	25
9. Universal Mounting Rails Installation Specifications	26
10. SCSI III Standards Overview	62
11. SCSI Differential Example	66
12. Piggy Back Connectors	69
13. Differential External Narrow Bus	79
14. Differential External Wide Bus	79
15. Comparison between SCSI and SSA Subsystem	82
16. IBM Expandable Storage Plus SCSI	85
17. IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer	87
18. Planar Block Diagram of the RS/6000 Model B50	92
19. PowerPC 604e Microprocessor Logical Block Diagram	97
20. System Management Services - Main Menu	119
21. SMS Submenu Multiboot	120
22. SMS Submenu Utilities	121
23. SMS Submenu - Firmware Passwords	122
24. Yellow Dog Linux Welcome Screen	160
25. Language Selection	161
26. Keyboard Type Selection - 1	161
27. Keyboard Type Selection - 2	162
28. Installation Method - 1	162
29. Installation Method - 2	163
30. Empty Keyboard Configuration File Error Message	163
31. Install/Upgrade Selection	164
32. Disk Partitioning	164
33. Filesystem Setup - 1	167
34. Filesystem Setup - 2	168
35. Swap Space Setting	168
36. Formatting Partitions	169
37. Selecting Packages	170
38. Installation Log	170
39. Mouse Setting - 1	171
40. Mouse Setting - 2	171

41. Network Configuration - 1	172
42. Network Configuration - 2	172
43. Network Configuration - 3	173
44. Timezone Configuration	173
45. Services Selection	174
46. Printer Configuration	174
47. Password Setting	175
48. Authentication Selection	175
49. Message About Bootloader	176
50. Monitor Selection/Setting	176
51. Resolution Selection for Display	177
52. Linux Installation Complete Message	177

Tables

1. SPEC and LINPACK Performance	9
2. Multiuser Performance	9
3. RS/6000 7046 Model B50 - System Board Components Location	13
4. Publications Shipped with the Model B50	17
5. Optional Memory for Model B50	21
6. List of Orderable Monitors for Model B50	22
7. TCP Streaming Tests	40
8. 16-bit SCSI Maximum Transfer Rates	44
9. Scorpion (DCHS) Adapter Rate for Model F50/H50	45
10. Sailfin (DGHS) Adapter Rate for Model F50/H50	45
11. Devices Location Codes	50
12. Communication, Graphics, and Storage Adapters Summary	56
13. Overview of SCSI-III Standards	63
14. Differential SCSI Modes	67
15. Single-Ended Ultra SCSI Adapter-to-First Device Cables	73
16. Device-to-Device Cables for Single-Ended Installations	74
17. 16-bit SCSI-II System-to-System Cable	75
18. Terminators for Single-Ended Installations	75
19. Differential-Ended Ultra SCSI Adapter-to-First Device Cables	76
20. Device-to-Device Cables for Differential-Ended Installations	76
21. Terminator for Description-Ended Installations	78
22. Cabling for the PCI Differential Ultra SCSI Adapter	79
23. Comparing SSA and SCSI	82
24. Internal Disk Drives for the Model B50	84
25. Internal CD-ROM Drives for Model B50	84
26. IBM 7133 Serial Disk System Advanced Models at a Glance	88
27. PowerPC and Bus Specification	95
28. Selected Service Provider Offerings	130
29. Service Providers Utility Pack Components	137
30. Features and Benefits of SecureWay Firewall	149
31. Test Machine Specifications	158

x RS/6000 7046 Model B50 Handbook

Preface

The RS/6000 7046 Model B50 is the first model in a new line of servers designed exclusively for the Internet Service Provider (ISP) and Application Service Provider (ASP).

The RS/6000 hardware Reliability, Availability, and Serviceability (RAS) design and the AIX UNIX leadership combine to provide a robust solution for an industry measured by availability. Aggressive pricing and an included AIX software license make the Model B50 the most affordable model in the RS/6000 product line. These attributes form the basis of a complete solution when the Model B50 is combined with IBM Web server or Firewall offerings.

With a slide out draw design, balanced mechanical chassis, ergonomically-located grab handles, and access to the adapters and disk drives requiring no tools, expansion is a simple task.

To help the Model B50 better fit in your DP strategy, it is one of the first Linux-enabled platforms in the RS/6000 product line.

This publication introduces the Model B50, discusses various configurations, and provides an overview of complete ISP and ASP software solutions.

The Team That Wrote This Redbook

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

The principle author of this publication is:

Tetsuroh Maeda IBM Japan

The project that produced this publication was managed by:

Scott Vetter IBM Austin

Thanks to the following people for their invaluable contributions to this project. Without their help, this publication would have been impossible:

Frank Hanson III IBM Austin

Earl Peterson IBM Austin

Richard Talbot IBM Austin

Ven Tenjarla IBM Austin

Comments Welcome

Your comments are important to us!

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

- Fax the evaluation form found in “IBM Redbooks Evaluation” on page 205 to the fax number shown on the form.
- Use the online evaluation form found at <http://www.redbooks.ibm.com/>
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. RS/6000 Introduction

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

This publication focuses on the IBM RS/6000 7046 Model B50 high-density rack-mount server (it needs only 2U per server). The Model B50 is one of IBM's first server offerings to specifically target the Internet and application server provider segments where small footprint and high-performance Web and application serving are critical business requirements and is focused on ISP and ASP environments where database transaction, e-commerce, and dynamic content demand this level of price and performance.

Figure 1 shows the Model B50 installed in an RS/6000 S00 rack. Versatile mounting rails enable the Model B50 to be mounted in almost any industry standard 19 inch rack.



Figure 1. RS/6000 7046 Model B50 Installed in S00 Rack

The RS/6000 naming convention is patterned similarly to many IBM product lines. A product is represented as a four-digit machine code followed by a model number. A unique machine code is assigned to each series of machines that fit together using similar designs or customer solutions. The first character of the model number groups machines with a specific enclosure together, and the remaining digits are typically used to indicate performance within similar models. This convention is not without variations.

The following sections provide an overview of the Model B50, the subject of this publication.

1.1 Basic ISP Server Topology

In Figure 2, the basic topology of an ISP environment is shown. A discussion of the featured components are discussed in the following.

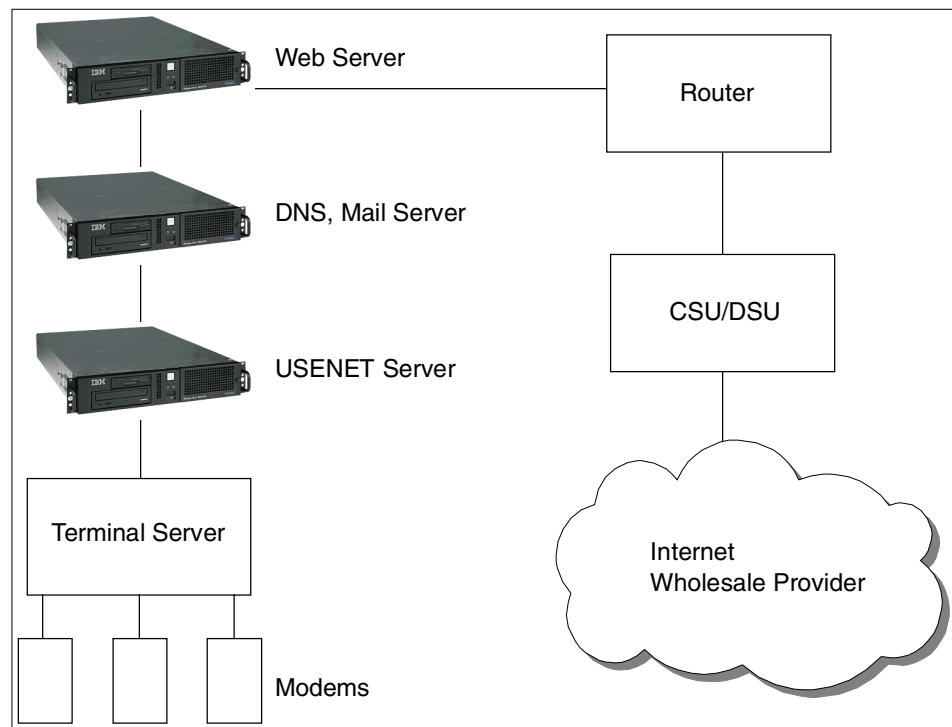


Figure 2. Typical ISP Topology

The basic configuration used in a startup ISP consists of the following:

- A series of servers providing the Internet services. Typically, Web server and USENET server functions are separated and scaled appropriately. A firewall server can be added between the router and the Internet.
- The Terminal Server is usually a dedicated unit that controls the attached modems and provides remote access for system administrators.
- The Router allows the various servers to work together on a subnet and appear, if required, as a single registered domain.
- The Channel Service Unit/Data Service Unit (CSU/DSU) provides the interface between the LAN data and the WAN connection provided by the wholesale provider. This function is sometimes integrated with the router.
- The wholesale provider may provide some of the functions, such as USENET, or a simple pipe to the Internet.

1.2 RS/6000 7046 Model B50 Overview

The IBM RS/6000 7046 Model B50 is a high-density, rack-mount server package offered at a competitive price. The Model B50 enables high-density packaging of a variety of applications for the Internet and corporate intranets with a very small footprint that can fit up to 20 servers per standard 19 inch telco rack.

The Model B50 is a uniprocessor system that provides enhanced performance by utilizing a 375 MHz PowerPC 604e processor and an enhanced memory controller. With this memory controller, the Model B50 uses SDRAM memory and an 83 MHz memory bus speed. The system memory can be expanded up to 1 GB.

With Ethernet and Ultra SCSI controllers integrated on the planar, two PCI slots, and four bays, the Model B50 is ready for expansion and growth.

The minimum Model B50 configuration includes a processor, CD-ROM, memory, and hard disk. The processor and CD-ROM will default. If you do not select memory or a hard disk, the system will default with 128 MB (# 4150) and a 9.1 GB disk drive (# 3112).

The Model B50 is shown in Figure 3.



Figure 3. RS/6000 7046 Model B50

The Model B50 supports a GXT130 2D graphics adapter, if a graphical desktop environment is required. In addition, a robust set of disk drive and communications features are available. The AIX Version 4.3 operating system for one to two users is included in the cost of the Model B50 and can be preinstalled if desired.

1.2.1 RS/6000 7046 Model B50 Key Attributes

The following is a list of the key attributes of the RS/6000 Model B50:

- Rack-mountable package.
- Permits installation of up to 1 GB of SDRAM memory in the server.
- Comes with Ethernet and Ultra SCSI controllers integrated on the planar.
- Enables high-density packaging of Internet and corporate intranet applications.
- Space-saving size provides packaging of up to 20 servers, 120 discrete network connections, and over 700 GB of online storage on a rack.
- Allows choice of exceptional AIX operating system 4.3.2 (by special request) or 4.3.3 with reliability, availability, and system-management features, or the Linux operating system (available from the Web).
- Ultra2 SCSI and SSA RAID storage external device options for high-performance access to high capacity, affordable disk storage.

The Model B50 is designed to help customers run many mission-critical business applications and networks twenty-four hours a day, seven days a week.

1.2.2 RS/6000 7046 Model B50 Feature Summary

The following is a summary of the key RS/6000 Model B50 features:

- Standard features include:
 - Microprocessor: 375 MHz PowerPC 604e processor
 - Level 2 (L2) cache: 1 MB
 - RAM (memory): 128 MB or 256 MB of ECC SDRAM memory
 - Memory slots: Four
 - Ports: One parallel and two serial
 - Internal disk drive: 9.1 GB or 18.2 GB Ultra SCSI disk drive
 - Disk/media bays: Four
 - Expansion slots: two PCI slots (one long and one half)
 - PCI bus speed: 33 MHz
 - 32X max speed CD-ROM
 - 1.44 MB 3.5-inch diskette drive
 - Integrated 10/100 Mbps Ethernet controller (IEEE 802.3 compliant)
 - Integrated Ultra SCSI controller
 - Tablet, keyboard, mouse, and audio ports
 - Tool-less access to all system components
 - Front accessible disk storage bays
- AIX operating system:
 - Version 4.3.2 (PRPQ) or Version 4.3.3
- Linux operating system:
 - Terra Soft Solutions' Linux 2.2 is available for the Model B50 by accessing: <http://www.yellowdoglinux.com/>
 - Support for Linux is not provided by IBM
- System expansion:
 - RAM: Up to 1 GB
 - Internal storage: 36.4 GB max
 - External disk storage: Maximum storage SCSI/SSA: Over 700 GB
 - A wide range of adapters available to support LANs and WANs.
- System dimensions:
 - 3.5" H x 17.6" W x 24.1" (standard 2U rack mount) minimum to 29.6" maximum D (88 mm x 447 mm x 612-752 mm)
 - Weight: 14.5 kg (32 lb) minimum to 15.9 kg (35 lb) maximum

- Operating environment
 - Temperature: 10° to 40°C (50° to 104°F)
 - Relative humidity: 8 to 80 percent
 - Operating voltage: 100 to 127 V AC or 200 to 240 V AC 50/60 Hz (manual switch)
- Warranty: On-site for one year

1.3 Performance Results

This section contains the expected values of several industry standard performance measurements for the RS/6000 Model B50.

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

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

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

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

1.3.1 SPEC95 Software Benchmark

SPEC95 is the forward step in the performance measurement of the core of the system. It covers the CPU, caches, memory, and compiler. The programs and data sets that make up the suite cannot load entirely into cache, making the benchmark more representative of real workloads. SPEC has also

standardized the compiler settings so that the results for *base* measurements are more comparable between suppliers.

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

SPEC95 contains two suites of benchmarks:

CINT95 Measures and compares computational-intensive integer performance

CFP95 Measures and compares computational-intensive floating point performance

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

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

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

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

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

There are several different ways to measure computer performance. One way is to measure how fast the computer completes a single task. This is a speed

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

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

SPECint95	SPEC component-level benchmark that measures integer performance. Result is the geometric mean of eight tests that comprise the CINT95 benchmark suite. All of these are written in C language.
SPECint_base95	The result of the same tests in CINT95 with a maximum of four compiler flags that must be used in all eight tests.
SPECint_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95).
SPECint_base_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95) with the restrictive compiler options.
SPECfp95	SPEC component-level benchmark that measures floating point performance. Result is the geometric mean of ten tests that comprise the CFP95 benchmark suite. All of these are written in FORTRAN.
SPECfp_base95	Result of the same tests in CFP95 with a maximum of four compiler flags that must be used in all ten tests.
SPECfp_rate95	Geometric average of the ten SPEC rates from SPEC floating point tests (CFP95).
SPECfp_base_rate95	Geometric average of the ten SPEC rates from the SPEC floating point tests (CFP95) with the restrictive compiler options.
SPECweb96	Maximum number of Hypertext Transfer Protocol (HTTP) operations per second achieved on the SPECweb96 benchmark without significant degradation of response time. The Web server software is ZEUS v1.1 from Zeus Technology Ltd.

LINPACK SP	Single precision, n=100 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS).
LINPACK DP	Double precision, n=100 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS).
LINPACK TPP	Toward Peak Performance, n=1000 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS). ESSL Version 3.1.1 was used in this test.

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

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

1.3.2 Performance Summary

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

Table 1. SPEC and LINPACK Performance

Model	CPU	Clock Rate (MHz)	L1 (KB)	L2 (MB)	SPEC int95	SPEC int_base95	SPEC fp95	SPEC fp_base95	LINPACK K DP	LINPACK SP	LINPACK TPP
B50	604e	375	32/32	1.0	15.1	14.5	10.1	9.76	64.8	151.2	255.7

Table 2. Multiuser Performance

Model	CPU	Clock Rate (MHz)	L1 Cache (KB)	L2 Cache (MB)	Relative OLTP Perf.	SPEC int_rate95	SPEC int_base_rate95	SPEC fp_rate95	SPEC fp_base_rate95
B50	604e	375	32/32	1.0	6.0	-	-	-	-

1.4 Industry Trends and Directions

There is rapid growth of UNIX-based hardware, software, and services requirements from the ISP and ASP marketplace. The market is already a 40 billion dollar global market and continues to expand.

Because they have to follow their customers requirements, such as more applications, better performance, and specialized services, IBM must continually introduce new solutions to meet our customers growing need.

IBM is continuously reviewing plans for these customers to expand their performance and capabilities. Future 7046 model offerings may include additional integrated interfaces, more disk storage, and faster I/O in a 2U format. Another possibility is to expand to a 4U SMP design, additional internal storage, and industrial standard NEBS compliance.

Chapter 2. Detailed Product Description

This chapter takes an in-depth look at the hardware packaging and features that comprise the IBM RS/6000 7046 Model B50 Server. This server uses the PowerPC architectures that offer improved memory throughput, better I/O performance, and flexible configurations.

2.1 IBM RS/6000 7046 Model B50 Product Description

The IBM RS/6000 7046 Model B50 is a powerful PowerPC 604e uniprocessor with a 32-bit implementation of the PowerPC architecture. The Model B50 is well-suited for customers like ISP or ASP whose services include Web applications, firewall, cache, file, e-mail, DNS, USENET, and so on.

Figure 4 shows a Model B50 with the access cover removed.




Figure 4. RS/6000 7046 Model B50

The RS/6000 7046 Model B50 is designed to allow for system growth in memory, disk, and media devices as the customer's needs grow. By combining AIX reliability and function with industry standard bus architectures, the Model B50 offers exceptional value for multiuser

applications and provides an excellent solution in the entry to the midrange server marketplace.

2.1.1 Standard Features

In this section, the standard features of the Model B50 and its most important characteristics are described.

2.1.1.1 Processor Subsystem

The PCI-based IBM RS/6000 Model B50 features the PowerPC 604e microprocessor with the internal clock speed of 375 MHz. With the internal processor clock to system bus clock frequency ratio set at 9:2, the system bus is run at 83 MHz. The processor 604e can support bus clock ratios of 1:1, 3:2, 2:1, 5:2, 3:1, 7:2, 4:1, 9:2, and 5:1.

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

The Model B50 L2 cache subsystem consists of a L2 cache controller and L2 memory. L2 cache memory consists of on-board Tag-Ram and 1 MB parity synchronous pipelined burst SRAM.

The Model B50 has two cache copy-back write buffers.

2.1.1.2 System Planar Board

One planar is installed in every 7046 Model B50. It provides the system bus and card connections for the riser card, a voltage regulator and clock module, four memory DIMM slots, power supply connections, I/O connections, and a pin-grid array socket for the PowerPC processor. Only single processor boards are available.

Jumpers are available to activate automatic power-on when utility power is present to enable automatic reboot after a power failure (J15) and jumpers are available to deactivate the external SCSI adapter to prevent unintentional data tampering (J35, J36).

Figure 5 on page 13 shows the system board location for a Model B50; Table 3 on page 13 lists the purpose of the various connections.




Figure 5. RS/6000 7046 Model B50 - System Board Locations

Table 3. RS/6000 7046 Model B50 - System Board Components Location

Location	Description	Location	Description
B1	Battery connector	J27, J30	System fan connector
J2	5x5 Auxiliary 5v connector	J28	Memory connector #2
J3	Audio input connector	J32	Privileged-access password jumper

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

2.1.1.3 Memory Subsystem

The Model B50 supports four industry standard DIMM sockets and can support up to 1 GB of system memory. Memory DIMM size is 128 MB or 256 MB. The memory DIMM should be installed from memory connector #4 then #3 to #1.

2.1.1.4 Bays

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

- 3.5-inch diskette drive bay

It accommodates the standard 1.44 MB diskette drive.

- 5.25-inch CD-ROM drive bay

It accommodates the standard 32X CD-ROM drive with the convenience of loading a CD without a caddy.

- 5.25-inch disk drive bay.

It accommodates the standard 9.1 GB Ultra SCSI disk drive that can be upgraded in the initial machine order to 18.2 GB Ultra SCSI disk. For more detailed information about the disk drives, see Table 24 on page 84.




Figure 6. RS/6000 7046 Model B50 - Front View

The empty bay (bay 3 or bay 4) can accommodate a 9.1 GB (# 3112) or 18.2 GB (# 3113) hard disk drive.

2.1.1.5 Internal CD-ROM Drive

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

For more detailed technical information about the CD-ROM drive, refer to Table 25 on page 84.

2.1.1.6 Standard I/O Ports

The system provides the following standard I/O ports:

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

2.1.1.7 PCI Adapter Slots

The Model B50 contains two 32-bit PCI slots. Following is a brief description according to their slots number:

- Slot 2 is PCI primary, long card available.
- Slot 1 is PCI secondary, short slot only.




Figure 7. RS/6000 7046 Model B50 - Rear View

2.1.1.8 Integrated Ethernet Adapter

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

2.1.1.9 Integrated Ultra SCSI Controller

The Model B50 has an integrated Ultra SCSI controller, which is operated at a speed of 40 MB/s.

2.1.1.10 AIX Operating System

The RS/6000 Model B50 includes a license for a full-function AIX Version 4.3 package at no additional charge.

2.1.1.11 Year2000 and EuroReady

The IBM RS/6000 7046 Model B50 is Year 2000 ready.

When used with AIX 4.3.2 or later, this product is EuroReady.

2.1.1.12 Publications

The following publications shown in Table 4 will be shipped or are available with the Model B50.

Table 4. Publications Shipped with the Model B50

Order Number	Title
SA38-0562 (shipped)	RS/6000 7046 B50 Setup Guide
SA38-0563 (shipped)	RS/6000 7046 B50 Users Guide
SA38-0554 ¹	RS/6000 7046 B50 Service Guide
SA23-2690 (shipped USA only)	Customer Support Information
SA23-2652 (shipped)	System Unit Safety Information
Customer Installable Options Library CD-ROM	
Warranty Booklet (USA customers only) ²	
Notes: ¹ The Service Guide is available, all other publications listed are shipped. ² The CD-ROM is not orderable, it is shipped only at manufacturing. No form number is available.	

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

2.1.2 Reliability, Availability, and Serviceability (RAS)

Following are the features that provide the IBM RS/6000 Model B50 reliability, availability, and serviceability.

2.1.2.1 Reliability, Fault Tolerance, and Data Integrity

The AIX serviceability features for error logging, analysis, and reporting are not available to those using the Linux operating system.

The reliability of the Model B50 system starts with reliable components, devices, and subsystems. During the design and the development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing, systems go through a testing process to ensure the highest product quality level.

The Model B50 system memory offers Error-Checking and Correcting (ECC) fault-tolerant features. ECC corrects environment-induced single-bit intermittent memory failures, as well as single hard failures. With ECC, the majority of memory failures will not impact system operation. ECC also provides double-bit memory error detection that protects data integrity in the event of the double-bit memory failures. The system bus and PCI buses are designed with parity error detection.

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

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

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

2.1.2.2 Fault Monitoring Functions

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

- Power-on-Self Test (POST) that checks the processor, L2 cache, memory, and associated hardware that are required for proper booting of the operating system every time the system is powered on. If a non-critical error is detected, or if the error(s) occur in the resources that can be removed from the system configuration, the booting process will proceed to completion. The error(s) are logged in the system Non Volatile RAM (NVRAM).
- Disk drive fault tracking is a facility that can alert the system administrator of an impending disk failure before it impacts customer operation.

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

2.1.2.3 Environment Monitoring Functions

Following is a list of the provided environment monitoring functions:

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

2.1.2.4 Error Handling and Reporting

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

2.1.2.5 Serviceability and Availability

The Model B50 is designed for customer setup of the machine and for subsequent addition of most features (adapters/devices). The Model B50 supports personnel to remotely log into a system to review error logs and perform remote maintenance.

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

Service mode allows checking of system devices and features.

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

Maintenance mode allows checking of most system resources.

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

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

2.1.2.6 Service Director for RS/6000

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

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

2.1.2.7 Online Customer Support (OCS)

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

2.1.3 Supported Optional Features

In this section, optional features are defined as internal optional devices, external adapters, and external subsystems that can be configured on the RS/6000 Model B50.

2.1.3.1 Memory

RS/6000 Model B50 uses ECC SDRAM memory. The Model B50 can be expanded up to 1 GB of memory. This allows the customer to easily upgrade

the memory cards by selecting one to four memory DIMMs (128 MB and 256 MB). The available memory specifications are provided in Table 5.

Table 5. Optional Memory for Model B50

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

Note

The DIMMs do not have to be installed in pairs.

2.1.3.2 Internal Disk Drives

Two internal Ultra SCSI disk drives, a 9.1 GB (# 3112) and 18.2 GB (# 3113), can be accommodated in the internal bays of the system. For more detailed information about the disk drives, see Table 24 on page 84.

2.1.3.3 Internal CD-ROM Drive

One 32X CD-ROM drive is in the system (# 2624) on media bay 2. You can find more detailed technical information about the CD-ROM drive in Table 25 on page 84.

2.1.3.4 Keyboard

The Stealth Black Quiet Touch Keyboards (# 87xx¹) are available in 38 languages. They have the same quiet and soft touch as the Pearl White Quiet Touch Keyboards (# 66xx¹). Like their predecessors, these keyboards have removable wrist rests for added comfort, and a 3-meter cable for added flexibility.

(¹ xx depends on the country.)

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

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

2.1.3.5 Mouse

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

2.1.3.6 Graphics Adapter

For 2D graphics applications, the Model B50 offers the POWER GXT130P accelerator card (# 2830). The GXT130P meets the graphics needs for many commercial applications and provides an entry point for workstation users. It is also a good price/performance choice for server environments.

Refer to 3.3, "Graphics Accelerators" on page 38 for more information.

2.1.3.7 Monitors

The IBM P76 and P260 color monitors are the premium graphics monitors for RS/6000 workstations. These monitors are for customers who demand image accuracy in a large screen format combined with performance, quality, reliability, and connectivity.

Table 6. List of Orderable Monitors for Model B50

Description	Feature Code	Screen Size (in inches)
P76 Color Monitor, Stealth Black	# 3627	17.0
P76 Color monitor, Pearl White	# 3629	17.0
P260 Color Monitor, Stealth Black	# 3628	21.0
P260 Color Monitor, Pearl White	# 3630	21.0

2.1.3.8 Software Preinstall

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

2.1.3.9 PCI Storage Adapters

The following PCI SCSI and SSA adapters are supported in the Model B50

- PCI Dual-Channel Ultra2 SCSI Adapter (# 6205)
- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI Differential Ultra SCSI Adapter (# 6207)
- PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)

- PCI Advanced SerialRAID Adapter (# 6225)

You can find more detailed information about the PCI storage adapters in 3.2, “PCI Storage Adapters” on page 35.

2.1.3.10 Communications Adapters

The following communications adapters are officially supported in Model B50.

- IBM High-Speed Token-Ring Adapter (# 4959)
- IBM PCI 10/100 Mbps Ethernet Adapter (# 2968)
- ATM Adapters
 - Turboways 155 UTP PCI ATM Adapter (# 2963)
 - Turboways 155 MMF PCI ATM Adapter (# 2988)
- FDDI Adapters
 - SysKonnnect SK-NET FDDI-LP SAS PCI Adapter (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS PCI Adapter (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS PCI Adapter (# 2743)
- Asynchronous Adapters
 - PCI 8-Port Asynchronous Adapter EIA-232/RS-422 (# 2943)
 - PCI 128-Port Asynchronous Controller (# 2944)
- 2-Port Multiprotocol PCI Adapter (# 2962)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)
- PCI IBM ARTIC960HX 4-Port Selectable Adapter (# 2947)

2.1.3.11 Supported Devices

These devices are supported through adapters supported above. Refer to the list of supported devices in the appropriate section of the AIX sales manual pages.

External Disk Subsystems

- 2104-DL1 Expandable Storage Plus Drawer
- 2104-TL1 Expandable Storage Plus Tower
- 7133-D40 SSA Disk Subsystem (Rack Mounted)

Tape Drive Subsystems

- 3490-F11 IBM Magnetic Tape Subsystem
- 3570-B00 IBM Magstar MP Tape Subsystem
- 3570-B01 IBM Magstar MP Tape Subsystem (Floor Library, 1 Drive)

- 3570-B02 Tape Subsystem (2 Drives)
- 3570-B11 IBM Magstar MP Tape Subsystem (Rack Mounted, 1 Drive)
- 3570-B12 IBM Magstar MP Tape Subsystem (Rack Mounted, 2 Drives)
- 3570-B1A IBM Magstar MP Tape Drive Model B1A
- 3570-C00 Magstar Magnetic Tape Subs * 1 Tape Lib (1), SCSI
- 3570-C01 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C02 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3570-C11 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C12 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3575-L18 IBM Magstar MP Tape Library DataServer - 180 Cartridge
- 3590-B11 High Performance Tape Subsystem
- 3590-C12 Silo Compatible Tape Subsystem
- 3590-E11 Magstar Tape Subsystem - Rack Mountable
- 7205-311 DLT Tape Drive Bridge Box
- 7206-005 4 GB 4 mm Tape Drive
- 7206-110 Cartridge Autoloader 4 mm, DDS-3, SCSI-2, DAT
- 7207-122 4 GB External SLR5 QIC Tape Drive
- 7207-315 13 GB 1/4" Tape Drive
- 7208-341 20 GB External 8 mm Tape Drive
- 7331-305 8 mm Tape Library
- 7337-305 Digital Linear Tape Subsystem
- 7337-306 Digital Linear Tape Subsystem

Printers, Plotters, and ASCII Terminals

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

Modems

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

2.1.3.12 Additional Universal Mounting Rails (# 6559)

These rails provide the customer the ability to mount additional equipment in IBM racks and most vendor racks. A single order of these rails will result in two rails (a pair) for a complete mounting solution.

Figure 8 shows how a single universal mounting rail is designed so it can adapt to a wide range of racks depths.




Figure 8. Universal Mounting Rails

Figure 9 on page 26 shows the recommended and minimum installation measurements required for proper Model B50 installation in an industry standard telco rack.

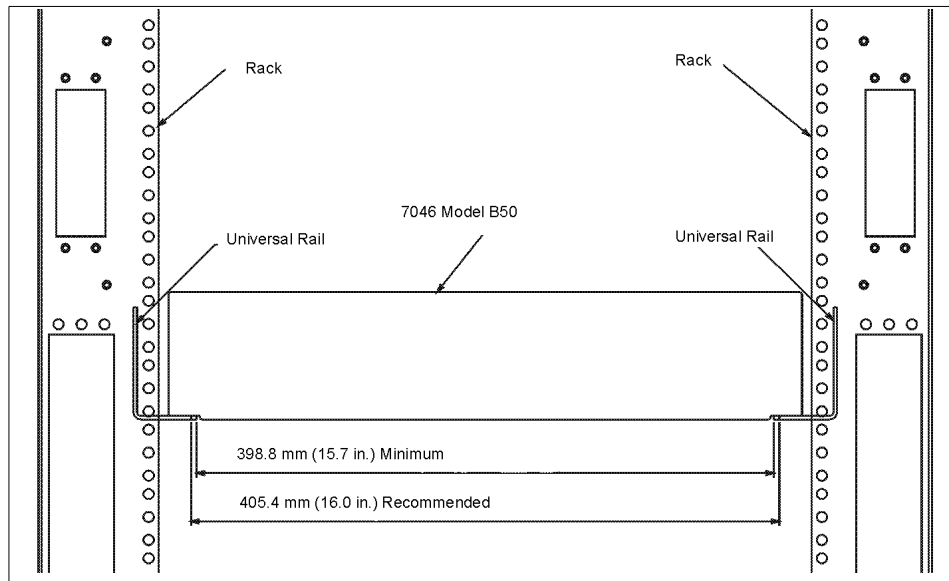


Figure 9. Universal Mounting Rails Installation Specifications

2.1.4 RS/6000 7046 Model B50 Adapter Placement Guidelines

This machine is designed for customers to update their configuration. When determining a slot in which to place a new adapter, you should start from the top of Table 12 on page 56. The slot number represents the order in which the slot should be used.

2.1.5 Service Package

The service package kit (# 3752) contains:

- Service publications
- Wrap plugs
- Diagnostic CD-ROM

The service publications include Maintenance Analysis Procedures (MAPs) that are not common to other RS/6000 systems. Also included is the *Diagnostics Information for Multiple Bus Systems*, SA38-0509, which contains information common to all multiple bus units and is used primarily for problem isolation by a customer's service specialist trained on the unit being serviced.

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

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

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

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

2.1.6 Options Library (Hard Copy)

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

Chapter 3. Adapters: Overview and Configuration

This chapter introduces the major communication and graphics adapters that are used with the RS/6000 7046 Model B50 rack mount server. Following the introduction, a discussion of PCI adapter configuration is given.

For a complete summary of the adapters supported on the Model B50, refer to Table 12 on page 56.

3.1 Communication Adapters

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

- Token-ring
- Ethernet
- ISDN
- FDDI
- ATM
- Multiple port
- WAN

3.1.1 Token-Ring Adapter

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

3.1.1.1 IBM High-Speed Token-Ring PCI Adapter (# 4959)

The High-Speed Token-Ring PCI Adapter provides 4/16 Mbps of token-ring performance in systems that incorporate the Peripheral Component Interconnect (PCI) bus interface. By providing 16 Mbps operation for both UTP Category-5 and STP type 1A cables, this adapter is capable of auto-sensing for both speeds (16 and 4 Mbps) and offers high performance without requiring significant cabling changes. The High-Speed Token-Ring PCI Adapter provides automatic ring speed selection and full duplex operation at both speeds (16 and 4 Mbps).

The following is a list of attributes for this card:

- Enhanced system interface for higher throughput and lower system utilization, coupled with RJ-45 and DB9 connectors

- Handles LAN environment requirements of today's servers and high-end workstations running I/O intensive applications on the network
- Operates at either 16 Mbps or 4 Mbps over Unshielded Twisted Pair (UTP) or Shielded Twisted Pair (STP)
- Automatic ring-speed detection
- External status LEDs providing visual indication of adapter and ring status
- Full-duplex LAN operation
- Supports system boot over token-ring

3.1.2 Ethernet Adapter

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

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

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

3.1.3 ISDN Adapter

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

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

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

Note

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

3.1.4 FDDI Adapters

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

Note

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

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

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

Features include:

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

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

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

Features include:

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

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

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

Features include:

- Occupies single PCI bus slot
- Compatible with the FDDI-ANSI X3T12 specification and FDDI Standard Series
- SMT 7.3 compliance
- AMD SUPERNET3 single chip FDDI design
- SysKonnnect-designed custom ASIC for buffer management with minimal CPU and bus load
- Four independent 128-byte FIFOs for buffering between PCI bus and buffer memory
- 128 KB static RAM on board
- Support under AIX SMIT and diagnostic routines
- Supports remote IPL and NIM with the appropriate system firmware

- Connects to Category-5 unshielded twisted pair FDDI concentrator ports using a shielded RJ-45 connector

Note

The AIX Device Driver diskette is shipped with the adapter.

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

3.1.5 ATM Adapters

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

3.1.5.1 PCI TURBOWAYS 155 UTP ATM Adapter (# 2963)

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

3.1.5.2 PCI TURBOWAYS 155 MMF ATM Adapter (# 2988)

The TURBOWAYS 155 PCI MMF ATM Adapter provides direct access to ATM networks. The TURBOWAYS 155 PCI MMF ATM Adapter provides dedicated 155 Mbps full-duplex connection using PVCs or SVCs and enables TCP/IP to run over an ATM network. The adapter also supports communication with devices located on an ATM network or bridged to a token ring, Ethernet, or other LAN.

Note

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

The TURBOWAYS 155 PCI MMF ATM Adapter is compatible with:

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

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

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

3.1.6 Multiple-Port Adapters

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

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

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

3.1.6.2 PCI 128-Port Asynchronous Controller (# 2944)

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

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

These RANs provide 16 EIA-232 (# 8137) or 16 RS-422 (# 8138) ports, yielding a total of 128-ports in a maximum configuration.

3.1.7 WAN Adapters

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

3.1.7.1 PCI 2-Port Multiprotocol Adapter (# 2962)

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

- X.21
- V.24
- V.35
- V.36

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

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

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

A RJ45 T1 cable is available through feature 2709 and a E1 cable through feature 2710.

3.2 PCI Storage Adapters

IBM RS/6000 Model B50 supports the following PCI-based storage SCSI and SSA adapters:

- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI Differential Ultra SCSI Adapter (# 6207)
- PCI Dual-Channel Ultra2 SCSI Adapter (# 6205)
- PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)
- PCI Advanced SerialRAID Adapter (# 6225)

3.2.1 SCSI Adapters

This section describes the available SCSI adapters for the Model B50.

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

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

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

3.2.1.2 PCI Differential Ultra SCSI Adapter (# 6207)

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

3.2.1.3 PCI Dual-Channel Ultra2 SCSI Adapter (# 6205)

The PCI Dual-Channel Ultra2 SCSI adapter enables you to use internal or external single-ended or Low-voltage Differential (LVD) Small Computer System Interface (SCSI) devices with computers containing a Peripheral Component Interconnect (PCI) type bus. This adapter provides two channels for SCSI devices (channel A and B channel). The PCI Dual-Channel Ultra2 SCSI Adapter has both an internal and a external connector on each channel. Only one connector (internal or external) can be used on each channel.

This adapter uses two Very High Density Cable Interconnect (VHDCI) 68-pin connectors allowing attachment to various LVD and SE devices. An adapter (# 2118) can be used to interface with 68-pin SE P-connectors.

3.2.1.4 PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)

The PCI 3-Channel Ultra2 SCSI RAID Adapter is a non-bootable high performance Ultra2 SCSI RAID adapter providing RAID 0, 1, or 5 capability and can address up to forty-five 16-bit SCSI-2 physical disk drives on three independent SCSI buses. To increase the data writing performance, a 32 MB fast-write cache is provided as a resident part of this adapter. The cache uses a nonvolatile RAM. The RAM can be removed from a failing adapter and installed in a replacement adapter and maintain cache integrity.

Two external ports provide connectivity to an IBM 2104-DL Expandable Storage Plus Drawer or a 2104-TL Expandable Storage Plus Tower at up to 80 MB/s SCSI data rate. The two external ports may also attach to a 7131-105, but speeds will be limited to 20 MB/s.

3.2.2 SSA Adapter

The PCI Advanced SerialRAID Adapter (# 6225) is supported in the Model B50.

3.2.2.1 PCI Advanced SerialRAID Adapter (# 6225)

The PCI Advanced SerialRAID adapter (# 6225) has four SSA connectors that are arranged in two pairs. Connectors A1 and A2 are one pair; connectors B1 and B2 are the other pair.

This adapter provides data transfer rates up to 160 MB/s per loop. It also provides eight-initiator, non-RAID capability, or two-initiator RAID capability. The adapter accepts a 32 MB SSA Fast-Write Cache Option Card (# 6235) that improves performance in both RAID 5 and non-RAID one-initiator applications.

The SSA links must be configured as loops. Each loop is connected to a pair of connectors at the SSA adapter card. These connectors must be a valid pair (that is, A1 and A2 or B1 and B2); otherwise, the disk drive modules on the loop are not fully configured, and the diagnostics fail. Operations to all the disk drive modules on a particular loop can continue if that loop breaks at any one point.

The Advanced SerialRAID Adapter when operated in a RAID 5 configuration will support (2+P) to (15+P) arrays and up to six (15+P) arrays. The adapter also supports Hot Spares in RAID 5 mode. The Advanced SerialRAID Adapter also supports connectivity to external disk enclosures and internal RS/6000 SSA configurations. Optional SSA Fiber-Optic Extender is also supported any supported RS/6000 system can be set up to boot from an Advanced SerialRAID Adapter (# 6225), 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 with AIX Version 4.3.3 or later software, native boot capability is supported as part of AIX Version 4.3.3 or later software.

3.3 Graphics Accelerators

This section discusses the GXT 130P graphics accelerator. Be sure to check the RS/6000 configurator for the application and availability of any RS/6000 adapter.

3.3.1 POWER GXT130P (# 2830)

The POWER GXT130P Graphics PCI Adapter is a high-performance PCI graphics adapter that accelerates and enhances your system unit video. POWER GXT130P Graphics PCI Adapter has no hardware switches to set. Mode selection is made via the software. Connection to the video monitor is made through a high density 15 pin D-shell connector.

- 8 MB SDRAM
- 32-bit PCI bus interface
- 8-bit color support
- Screen resolutions:

640x400 at 60 - 85 Hz vertical refresh
800x600 at 60 - 85 Hz vertical refresh
1024x768 at 60 - 85 Hz vertical refresh
1280x1024 at 60 - 85 Hz vertical refresh
1600x1200 at 75 - 85 Hz vertical refresh

The POWER GXT130P is a versatile, low-priced graphics accelerator for the RS/6000 B50 system. It accelerates X Windows graphics applications, such as displaying business graphics, data base status, CASE tools, and activity for securities trading. The GXT130P can support a network control display for monitoring activity and assisting in load balancing, performance analysis, and capacity planning.

The GXT130P is also ideal for process control, where, in selected configurations, up to four monitors can be operated concurrently from the same RS/6000 system for displaying graphics and text for dozens of process points.

The GXT130P supports screen resolutions of 640x480, 800x600, 1024x768, and 1280x1024 all at 8-bits per pixels. It satisfies the ISO 9241 standard of visual quality at all these resolutions, refreshing the screen at up to 85 Hz for all resolutions, including the 1280x1024 8-bit resolution.

The GXT130P graphics accelerator is a half-length PCI card with a standard 15-pin DIN output connector.

3.4 Performance Notes for PCI Communications and Storage Adapters

The purpose of this section is to provide some *rules of thumb* for the performance of various network and disk adapters.

Throughput and CPU utilization for a few typical systems is provided. This information can then be used to approximate throughput and CPU utilization with multiple adapters and for other systems.

Disk throughput is discussed as well as how the disk throughput may be the bottleneck when using higher speed LAN adapters.

3.4.1 Adapter Placement

Please consult 3.6, “Summary Chart of Adapters, Size, and Placement” on page 56, for slot placement guidelines and limitations that may exist on the number of adapters that can be supported for connectivity and the number that can be supported for maximum performance.

Several PCI machines have secondary PCI buses bridged onto a primary PCI bus. Some medium to high-speed adapters perform slower on these secondary bus slots and some adapters are recommended to not be used in these slots. Machines with some secondary PCI slots include E30, F40, and SP 332 MHz SMP wide nodes.

3.4.2 Networking Throughput

Table 7 on page 40 details the TCP streaming tests for various communications adapters. Adapter throughput is shown together with CPU utilization of a 233 MHz RS/6000 E30 and a 333 MHz 4-way RS/6000 F50. These values are provided *as is*, and were recorded on development level machines. Actual performance may vary considerably.

Results provided here are for *user payload data* that a user program would see. The actual bit rate on the wire would be higher due to the added TCP/IP protocol headers and data link headers.

Results are for sending and receiving packet sizes in the range of 8 KB to 64 KB. Smaller size I/Os will result in higher CPU utilization and lower throughput. For reference, AIX FTP uses 64 KB send and receive buffer size, but small files may result in smaller I/Os. Peak performance is normally found for 16 KB read and write sizes, but will be similar across the 8 KB to 64 KB range.

CPU utilization is the CPU utilization taken from the busiest of the sending or receiving systems (the receiver usually uses somewhat more CPU).

These results were taken on otherwise idle machines and networks.

The data is transferred between two application programs (Netperf) using the sockets API using the TCP/IP protocol. The transfers are memory to memory, so they do not account for any disk I/O or file system overhead.

Therefore, these numbers represent *best case* performance and should be adjusted downward to compensate for real-world issues, such as smaller files, other network and system activity, and network application overhead.

Table 7. TCP Streaming Tests

IBM PCI Adapter	Workload	Mbps ¹	MB/s ²	E30 233 MHz	F50 332 MHz 4-way
Ethernet in 10 Mbps HDX ³	1-way TCP stream ⁴	9.0	1.1	9%	2.5%
Ethernet in 10 Mbps FDX ⁵	1-way TCP stream	9.0	1.1	9%	---
Ethernet in 10 Mbps FDX	2-way TCP stream ⁶	18.0	2.2	18%	4.3%
Ethernet in 100 Mbps HDX	1-way TCP stream	84.0	10.3	---	22%
Ethernet in 100 Mbps FDX	1-way TCP stream	94.9	11.5	100% ⁷	25%
	2-way TCP stream	102.0	12.5	100% ⁷	---
	2-way TCP stream	176.0	21.5	---	34%
ATM 155, Ethernet LANE ⁸ MTU ⁹ 1500 ¹⁰	1-way TCP stream	121.0	14.7	100% ⁷	---
	1-way TCP stream	124.0	15.0	---	33%
	2-way TCP stream	121.0	14.7	100% ⁷	---
	2-way TCP stream	180.0	22.0	---	40%
ATM 155, default MTU 9180	1-way TCP stream	129.0	15.8	60%	---
	1-way TCP stream	133.0	16.2	---	14%
	2-way TCP stream	165.0	20.3	87%	---
	2-way TCP stream	222.0	27.1	---	20%
FDDI, default MTU 4350	1-way TCP stream	86.0	10.5	62%	14%
Token-Ring, 16 Mbps MTU 1492	1-way TCP stream	14.8	1.8	16%	4.5%
Token-Ring, 16 Mbps MTU 8500	1-way TCP stream	15.8	1.9	8%	1.8%
Token-Ring, 16 Mbps MTU 1492, FDX ¹¹	2-way TCP stream	29.2	3.6	32%	8%

IBM PCI Adapter	Workload	Mbps ¹	MB/s ²	E30 233 MHz	F50 332 MHz 4-way
Token-Ring, 16 Mbps MTU 4000, FDX	2-way TCP stream	29.2	3.6	15%	4%
Gigabit Ethernet, default MTU ^{1,2} 1500	1-way TCP stream 1 TCP session	262	32.0	---	42%
	1-way TCP stream 2 TCP sessions	275	33.5	---	56%
	1-way TCP stream 8 TCP sessions	277	33.8	---	60%
	2-way TCP stream 2 sessions	345	42.1	---	65%
"Jumbo Frame" MTU 9000	1-way TCP stream 1 TCP session	328	40.1	---	30%
	1-way TCP stream 2 TCP sessions	475	58.0	---	47%
	1-way TCP stream 8 TCP sessions	542	66.1	---	57%
	2-way TCP stream 2 TCP sessions	436	53.2	---	40%
	2-way TCP stream 4 TCP sessions	577	70.4	---	58%

Notes:

¹ Mbps is computed using 10**6 bits/second.

² MB/s is computed using 2**20 Bytes/second.

³ HDX - Half duplex.

⁴ TCP one-way stream program on machine A is doing *sends* to send data to program on machine B that is doing *receives*. This is similar to what FTP would do, but without the file I/O.

⁵ FDX - Full duplex.

⁶ TCP two-way stream has two pairs of programs, A sending to B and C sending to D. This places a full-duplex workload on the transport system and for adapters that have full-duplex media, the throughput should be about double the one-way streaming results, assuming you have enough CPU and I/O bus resource.

⁷ Limited by CPU speed.

⁸ LANE - LAN Emulation.

⁹ MTU - Maximum Transmission Unit.

¹⁰ ATM LAN Emulation (LANE) typically uses 1500 byte MTU size, so this will cause increased CPU utilization and lower performance than running classical IP with the larger 9180 byte MTU size.

¹¹ Full duplex token-ring requires a switch supporting full duplex.

¹² Default Maximum Transmission Unit (MTU) (packet size on the wire) are: Ethernet 1500 (maximum is 1500, except for Gigabit Ethernet where jumbo frames use MTU of 9000), token-ring 1492 (maximum is 17792), FDDI 4352 (maximum is 4352), ATM 9180 (maximum is 65527).

Larger MTU sizes result in lower system utilization for the same throughput because the system will have less packets and interrupts to handle. Maximum MTU sizes may be limited by network switches and routers. Running with MTU sizes larger than 16,384 may require use of RFC1323 option to allow use of larger TCP send and receive space greater than the 65,536 to have good performance. Without use of RFC1323, the MTU size should be limited to 16,384 bytes.

Note

Throughput will usually scale up linearly as multiple adapters are added until you reach about 90 percent CPU utilization or until the PCI bus throughput limit is reached. For example, on a Model F50, one 100 Mbps adapter in FDX mode will run 94.9 Mbps and use 25 percent CPU. This machine will run four adapters at 377 Mbps and use 90 percent CPU. Adding another adapter will only give 43 more Mbps, and the machine is then at 100 percent CPU.

If the machine is only a 2-way instead of a 4-way, double the CPU utilization that is shown in the table.

You should normally spread the adapters over multiple buses if the machine has multiple buses.

3.4.2.1 Discussion of Ethernet

Ethernet started life as a coax-based network that was CSMA/CD based and was, therefore, half duplex. This was termed 10base5, for 10 Mbps baseband and 500 meter coax segments. The CSMA/CD access would only allow the network to perform well if it was below 70 percent utilization because the collisions and backoff-retries became excessive. Ethernet has evolved rapidly and now supports 10baseT and 100baseT. The T means twisted pair that uses the RJ45 connector. 100 means 100 Mbps. The 10baseT and 100baseT can also support full-duplex so that the adapter can send and receive at the same time. However, the proper network infrastructure is required to support full-duplex operation.

The network interconnect devices (transceivers, hubs, switches) have also evolved. Network hubs only support half-duplex mode, but some may support one port with full duplex mode. To enjoy maximum performance, you need to use an Ethernet switch. A switch can switch packets between ports so that multiple concurrent 10 or 100 Mbps sessions are running over the network. For example, you can have machine A sending to machine B at 100 Mbps and machine C sending to machine D at 100 Mbps concurrently. This would not be possible with a lower priced hub. If two clients are sending to the same server, the server will still only be able to receive at the speed of the link (100 Mbps). However, the server could be sending to a third client at the same time it is receiving from one of the other clients. Most modern adapters and software support both half-duplex and full-duplex modes. The highest performance will be in full-duplex mode. However, you must be sure that the network switches in the network also support full-duplex if you use this mode. Some switches require that you manually configure for full-duplex mode.

Most modern switches and adapters will auto-configure and select the fastest mode. It is always good to check, however, because one adapter in half-duplex mode and another in full-duplex mode can cause performance problems. The IBM 10/100 adapter will show the negotiated mode by using the `entstat` command, for example `entstat -d en0`. The result can be seen at the bottom of the report.

3.4.2.2 Discussion on Token Ring

Token ring has also evolved and can support full-duplex now. This requires a special switch that can support this operation. An example is the IBM 8272 Token Ring switch. Some older PCI token-ring adapters may require a firmware update to support full-duplex mode.

3.4.2.3 Discussion on ATM

ATM is a circuit-switched network, much like the telephone system, where you have to establish a connection between two endpoints before you can communicate. Networks like Ethernet, token-ring, and FDDI are *broadcast*-based networks, where any station on the network can talk to any other station on the network without setting up a connection at the data-link level. With ATM, you must establish a connection at the data-link level.

ATM can support both classical IP and LAN Emulation (LANE). With classical IP, an administrator must configure the machine with either switched or permanent virtual circuits between endpoints. From that point on, all the normal TCP applications will work transparently. LAN Emulation was provided to make ATM look like a broadcast network to ensure that various protocols (like NetWare, for example) could run on top of ATM in a transparent way. LANE can configure Ethernet or token-ring interfaces on top of ATM. These interfaces will have the MTU size and bit rate set to cause ATM to emulate Ethernet or token-ring.

LANE should be avoided if at all possible because it forces ATM to use smaller MTU size, and this results in higher CPU utilization on the system; thus, the system can drive fewer adapters or do less other work.

You can configure both classical IP and LANE to run concurrently over the same ATM device. For example, LANE could run to some PC workstations, and classical IP could be used to back up the system with ADSM, so the backup can take advantage of the larger packets and higher performance.

One other factor to be aware of is that a 155 Mbps OC-3 ATM link will only provide about 134 Mbps of throughput in each direction. This is because the physical layer signaling overhead reduces the effective bandwidth in addition

to the 5 bytes of ATM cell header for each 53-byte ARM cell which consumes another nine percent overhead.

3.4.3 Disk I/O Throughput

For the faster network adapters, when transferring data to/from a file system, the disk may be the bottleneck. For example, if an SCSI disk can transfer data at 7 MB/s, that would be the bottleneck when using ATM that can transfer data at 16 MB/s when using Classical IP or when using 100 Mbps Ethernet, which can transfer data at 11 MB/s. Disk striping can be used to increase the transfer rate for sequential files. For example, two 7 MB/s disks striped would yield 14 MB/s, which would be close to the ATM limit.

SCSI has evolved rapidly, and the transfer rate of the adapter will depend on the specific SCSI technology used.

Some of the maximum possible transfer rates for SCSI are noted in Table 8. Various adapter and system implementation issues may limit the maximum rate that can be achieved.

Table 8 provides 16-bit SCSI maximum transfer rates.

Table 8. 16-bit SCSI Maximum Transfer Rates

SCSI Adapter	Maximum Transfer Rate (MB/s)
Fast/Wide	20
Ultra/Wide	40
Ultra2/Wide	80

Actual disk transfer rates will vary by the disk type, adapter type, system bus limits, and CPU speed. Typical transfer rates to a single disk will be in the range of 7 to 14 MB/s. Disk reads will usually be a little faster than writes. Disk transfer speeds and access times continue to improve over time. It is very difficult to tie the disk type to an IBM feature code because these can change over time. In the following tables, a disk type is listed (for example, DCHS), and this can be used to determine the exact type of disk an actual machine has. If you are ordering a new system, it will generally have the latest (fastest) technology disk.

The following table shows some transfer rates for sequential reads and writes to JFS files. The file size is 16 MB, and the read and writes are 32 KB bytes. Throughput will be similar if the read and write size is in the range of 4 KB to 256 KB, but CPU usage may vary. AIX default read-ahead is 32 KB. Applications that use raw I/O will result in much lower CPU usage because

raw I/O bypasses the JFS file system. Applications like FTP will only be able to read or write to a single file (disk) at a time, unless disk striping is being used. Disk striping is the only way to increase the throughput rate for applications that read or write sequentially to a single file.

Table 9 provides a table of the Ultra SCSI (SE or differential) and the DCHS rates on an Model F50/H50 332 MHz 4-way:

Table 9. Scorpion (DCHS) Adapter Rate for Model F50/H50

# of Disks	# of Adapters	JFS Read Seq. (MB/s @CPU%)	JFS Write Seq. (MB/s @CPU%)
1	1	9.4@7%	7.4@4%
2	1	18.6@12%	14.5@9%
4	1	25.0@16%	24.1@14%
8	2	50.0@35%	74.8@58%
12	3	74.8@58%	70.8@50%

Table 10 provides a table of the Ultra SCSI (SE or differential) and the 9.1 GB (#2913, DGHS) rates on an Model F50/H50 332 MHz 4-way:

Table 10. Sailfin (DGHS) Adapter Rate for Model F50/H50

# of Disks	# of Adapters	JFS Read Seq. (MB/s @CPU%)	JFS Write Seq. (MB/s @CPU%)
1	1	14.3@9%	13.6@8%
2	1	29.4@19%	21.8@13%
3	1	25.6@17%	25.5@15%
4	2	30.1@20%	27.9@17%

To find out what kind of disk you have, substitute the disk number for the x in the command given:

SCSI Disks:

```
lscfg -vl hdiskX | grep [machine type]
```

DPES 1080 MB 3.5" SCSI Disk (Withdrawn)

DORS (# 3905) 2.1 GB Ultra SCSI Disk (Withdrawn)

DDRS (# 3028) 4.5 GB Ultra SCSI Disk

DFHS (# 3035, 3037, 3034) 1.1, 2.2, 4.5 GB F/W SCSI Disk (1.7 inch)

DCHS (# 3081, 3090) 4.5 GB, 9.1 GB (1.7 inch)
DGHS (# 3029) 9.1 GB (also 18 GB 1.7 inch not shipped by RS6K)
DRVS (10,000 RPM drive, like DGHS but faster random I/O)

SSA Disks:

```
lscfg -vl pdiskX | grep [machine type]
```

DFHC	SSA 4.5 GB
DCHC	SSA 9.1 GB 1.6"
DGHC	SSA Marlin

3.4.3.1 Combined Network and Disk Performance

Real-world applications are a combination of disk and network I/O. Examples include using FTP to transfer files or doing ADSM backups over the network.

The preceding tables give some guidelines on the throughput and CPU needed to do the network and disk I/O. The CPU utilized to drive the network and the disk must be combined to find the total CPU required. What is unknown is the additional CPU cycles needed by the application. If the CPU is close to 90 percent or more, any cycles needed by the application will reduce the total throughput.

If the disk is slower than the network, the disk will be the bottleneck. Disk striping can help remove this bottleneck, or choose one of the faster disk adapters and/or disk drives.

If the network is slower than the disk, the network is the bottleneck. Moving up from 10 Mbps Ethernet to 100 Mbps Ethernet, for example, may be the answer. Using full-duplex media and switches (instead of hubs) will also help. Use of classical IP instead of LANE when using ATM can also remove the CPU or adapter as the bottleneck.

In some cases, the system CPU will be the bottleneck. If the CPU is above 90 percent utilized, then moving up to a faster processor or adding more processors should help. More CPUs may not help if you are limited to a single adapter. Avoid using small packet size adapters. Do not use ATM LANE if you can use classical IP. Use Jumbo Frames on Gigabit Ethernet if possible. These suggestions will reduce the CPU utilization, which may remove the CPU as being the bottleneck. If the CPU is the bottleneck, the throughput rate will be reduced because of the CPU cycles needed for the disk I/O and for the application.

FTP can transfer in both binary mode and in ASCII mode (default). Binary mode should be used whenever possible because it is faster. ASCII mode must scan the data, and this uses more CPU and will reduce the throughput rate if the system is CPU limited. Transferring directories of small files can reduce throughput. Combining small files into a single large file, for example, with the `tar` command, and then compressing the file can reduce the need for network bandwidth.

Some FTP tests were conducted on a 43P Model 150, which is a 375 MHz machine over various networks. The disk used was the DCHS Scorpion 4.5 MB drive.

Transferring a 100 MB file over ATM 155 Mbps in classical IP mode, a binary transfer ran at the disk speed of 9.3 MB/s and used 48 percent of the CPU. The transfer was disk bound because it is known that the Scorpion disk will transfer at 9.4 MB/s, per the preceding F50 tables, and use seven percent of the CPU on a 4-way 333 MHz machine. Thus, in round numbers, it should take about four times that much CPU on a uniprocessor of about the same speed. Thus, the 9.3 MB/s disk I/O should take about 28 percent of the system.

The ATM can transfer data at about 16 MB/s. From the E30 network table, it was determined that the E30 233 MHz uses 60 percent CPU to give 15.8 MB/s. Adjusting this CPU for the 375 MHz Model 150, you get $60 * 233/375 = 37$, or the Model 150 should use 37 percent CPU. (It is known from the actual test that it used 39 percent for the Netperf test, so our estimate is very close.) However, it only ran at 9.3 MB/s; so you need to reduce the network CPU part by $9.3/15.8 = .58$. $37 \text{ percent} * .58 = 21.5 \text{ percent}$. So the combined disk and network CPU usage would be 28 percent for the disk I/O and 21.5 percent for the network I/O for a total of 49.5 percent. It was measured at 48 percent. This example shows that these rules-of-thumb can be used to estimate performance.

The same FTP transfer when done with the file residing in memory (for example, the second transfer) ran at 16.2 Mbps, which is at the network media speed. The measured CPU usage averaged 59 percent on the remote machine (but peaked at 83 percent) and 53 percent on the local machine. This average rate appears to be close to the combined 27 percent for the disk transfer and 39 percent for the full speed media transfer ($27 + 39 = 66$ percent) that was predicted.

In this example, the FTP binary mode adds little CPU overhead above what you need for disk and network I/O. Other applications, like ADSM, can add

quite a bit of overhead to cover while working with small files and its accounting chores.

For these reasons, these numbers should be considered best case and should be adjusted to be conservative for real-world use.

3.5 Adapter and Device Configuration on PCI-Based RS/6000 Systems

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

3.5.1 Device Types

PCI-based RS/6000 B50 system supports the attachment of devices using the following buses:

- SCSI
- PCI
- ISA (limited to built-in support)

Devices on PCI-based RS/6000 systems can be of two types: native devices, such as the integrated SCSI controller and the serial controller, or attached (non-native) devices, such as the supported PCI adapters.

The device information required to configure the integrated devices is contained in the firmware and is passed to the operating system through the residual data written to NVRAM. Thus, integrated devices are configured automatically to AIX after system reboot.

Attached PCI devices and adapters can be configured automatically because the PCI Local Bus Specification defines a configuration method for PCI components. Provided that the device support software is installed, PCI and SCSI devices are configured automatically whenever the configuration manager program, `cfgmgr`, is run at system boot or run time, and when no conflict (for example, the same SCSI ID for two different SCSI devices) is found.

The ISA/EISA bus has no standard method of identifying adapters or their configuration requirements.

3.5.1.1 Device Configuration Database

Device information is contained in a predefined database or in a customized database that makes up the Device Configuration Database managed by the Object Data Manager (ODM).

The predefined database contains configuration data for all possible devices configurable for the system.

The customized database contains configuration data for all currently defined devices in the system.

The device information stored in the Device Configuration Database allows the automatic configuration of PCI devices on PCI-based RS/6000 systems whenever the Configuration Manager (`cfmgmr`) program is run at system boot and run time.

The Device Configuration Database that is currently used by the system is located in the directory that is specified by the ODMDIR environment variable. This is normally the `/etc/objrepos` directory.

In AIX Version 3.2, all device software was installed when the operating system was installed. In AIX Version 4, only those software packages for devices that were detected at the installation time are automatically installed.

Refer to the AIX Version 4 product documentation for more information on device management and enhancements to the Device Configuration Database.

3.5.2 Device Location Codes

The RS/6000 7046 Model B50 uses physical location codes in conjunction with AIX location codes to provide mapping of the failing field replaceable units. The location codes are produced by the system unit's firmware and AIX.

3.5.2.1 Physical Location Codes

Physical location codes provide a mapping of logical functions in a platform (or expansion sites for logical functions, such as connectors or ports) to their specific locations within the physical structure of the platform.

Location Code Format

The format for the location code is a string of alphanumeric characters separated by a dash (-), slash (/), pound sign (#), or period (.):

- The base location is all of the information before the slash (/) or pound sign (#). It identifies a device that is connected or plugged into the parent.

- Extended location information follows the slash (/). It identifies a device that is part of the parent, a connector, or a cable.
- Cable information follows the pound sign (#). It identifies a cable that is connected to the parent.

The following are examples:

- P1-C1 Identifies a CPU card C1 plugged into planar P1.
- P1-M1 Identifies a memory card M1 plugged into planar P1.
- P1-K1 Identifies a keyboard attached to connector K1 on planar P1.
- P1/S1 Identifies serial port 1 controller on planar P1, the connector for serial port 1, or the cable attached to the connector for serial port 1.
- P1-I2/E3 Identifies Ethernet controller 3 on the card plugged into slot 2 (I2) on planar P1, the connector for Ethernet controller 3, or the cable attached to Ethernet controller 3.
- P1-I2#E3 Identifies the cable attached to Ethernet controller 3 plugged into slot 2 (I2) on planar P1.

Sublocations, such as memory modules on a base memory card or a specific SCSI address, are identified by a period (.). The following are examples:

- P1-M1.4 Identifies DIMM 4 on memory card 1 on planar 1.
- 1-C1.1 Identifies CPU 1 on CPU card 1 on planar 1.
- P2/Z1-A3.1 Identifies a SCSI device with a SCSI address of LUN 1 at SCSI ID 3 attached to SCSI bus 1 from planar 2.
- P1-I2#E3.2 Identifies the second cable in a series of cables attached to Ethernet controller 3 in slot 2 (I2) on planar 1.

Depending on the AIX and firmware levels, AIX Diagnostics may include extended location information when identifying a planar or card. The extended location information is always included when identifying a cable or connector. Location codes with extended location information that display without a description identifying the devices always identify the cable attached to the port. See Table 11.

Table 11. Devices Location Codes

Name	AIX Location Code	Location Code	Physical Connection	Logical Identification
System Planar	00-00	P1		
Riser Card		P1.1	J39	

Name	AIX Location Code	Location Code	Physical Connection	Logical Identification
Memory DIMM 1	00-00	P1-M1	J11	
Memory DIMM 2	00-00	P1-M2	J28	
Memory DIMM 3	00-00	P1-M3	J19	
Memory DIMM 4	00-00	P1-M4	J39	
Diskette Drive	01-D1-00-00	P1-D1	J16	Base Address 0x03F0
Keyboard	01-K1-00-00	P1-K1	J59	Base Address 0x0060
Mouse	01-K1-01-00	P1-O1	J56	Base Address 0x0060
Audio Port	01-Q2	P1/Q2		
Diskette Port	01-D1	P1/D1	J16	Base Address 0x03F0
Keyboard Port	01-K1-00	P1/K1	J59	Base Address 0x0060
Mouse Port	01-K1-01	P1/O1	J56	Base Address 0x0060
Serial Port 1	01-S1	P1/S1	J24	Base Address 0x03F8
Serial Port 2	01-S2	P1/S2	J23	Base Address 0x02F8
Parallel Port	01-R1	P1/R1	J49	Base Address 0x0378
Ethernet Port	10-60	P1/E1	J45	Host Bridge ID00,Device 12
SCSI Port	10-80	P1/Z1	J33/J38	Host Bridge ID00,Device 16
Tablet Port	01-Q1	P1/Q1		
Card in PCI Slot 2	10-90	P1.1-I2	J3 on Riser Card (Top Connector)	PCI Bridge ID01, Device 18

Name	AIX Location Code	Location Code	Physical Connection	Logical Identification
Card in PCI Slot 1	10-b0	P1.1-11	J2 on Riser Card (Bottom Connector)	Host Bridge ID00, Device 22
CD-ROM		P1-Z1-A3		Primary SCSI Bus ID 3
Upper DASD		P1-Z1-A2		Primary SCSI Bus ID 2
Lower DASD		P1-Z1-A4		Primary SCSI Bus ID 4
Fan # 1		F1	J27/J30	
Operator Panel		L1	J41	
Power Supply		V1	J7/J8/J2/J9	

3.5.2.2 AIX Location Codes

The basic formats of the AIX location codes are:

For non-SCSI devices/drives AB-CD-EF-GH

For SCSI devices/drives AB-CD-EF-G,H

For planars, cards, and non-SCSI devices, the location code is defined as:

AB-CD-EF-GH

```

| | | |
| | | Device/FRU/Port ID
| | Connector ID

```

| devfunc Number, Adapter Number or Physical Location
Bus Type or PCI Parent Bus

- The AB value identifies a bus type or PCI parent bus as assigned by the firmware.
- The CD value identifies adapter number, adapter's devfunc number, or physical location. The devfunc number is defined as the PCI device number times 8, plus the function number.
- The EF value identifies a connector.
- The GH value identifies a port, address, device, or FRU.

Adapters and cards are identified with just AB-CD.

The possible values for AB are:

00	Processor bus.
01	ISA bus.
02	EISA bus.
03	MCA bus.
04	PCI bus used in the case where the PCI bus cannot be identified.
05	PCMCIA buses.
xy	For PCI adapters where x is equal to or greater than 1. The x and y are characters in the range of 0-9, A-H, J-N, P-Z (O, I, and lower case are omitted) and are equal to the parent bus's ibm, aix-loc Open Firmware Property.

The possible values for CD depend on the adapter/card:

- For pluggable PCI adapters/cards, CD is the device's devfunc number (PCI device number times 8, plus the function number). The C and D are characters in the range of 0-9 and A-F (hex numbers). This allows the location code to uniquely identify multiple adapters on individual PCI cards.
- For pluggable ISA adapters, CD is equal to the order the ISA cards defined/configured either by SMIT or the ISA Adapter Configuration Service Aid.
- For integrated ISA adapters, CD is equal to a unique code identifying the ISA adapter. In most cases, this is equal to the adapter's physical location code. In cases where a physical location code is not available, CD is FF.

EF is the connector ID. It is used to identify the adapter's connector that a resource is attached to.

GH is used to identify a port, device, or FRU. For example:

- For async devices, GH defines the port on the fanout box. The values are 00 to 15.
- For a diskette drive, H defines which diskette drive 1 or 2. G is always 0.
- For all other devices, GH is equal to 00.

For an integrated adapter, EF-GH is the same as the definition for a pluggable adapter. For example, the location code for a diskette drive is 01-D1-00-00. A second diskette drive is 01-D1-00-01.

For SCSI, the location code is defined as:

```
AB-CD-EF-G,H
| | | |
| | | | Logical Unit address of the SCSI Device
| | | Control Unit Address of the SCSI Device
| | Connector ID
| devfunc Number, Adapter Number or Physical Location
Bus Type or PCI Parent Bus
```

Where:

- AB-CD-EF are the same as non-SCSI devices.
- G defines the control unit address of the device. Values of 0 to 15 are valid.
- H defines the logical unit address of the device. Values of 0 to 255 are valid.

There is also a bus location code that is generated as 00-XXXXXXXX where XXXXXXXX is equivalent to the node's unit address.

3.5.2.3 Configuration Information

Information about the PCI adapters can be obtained by running the `lscfg` command.

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

```
lsattr -El mem0
```

This will show the size of memory in your machine and how much usable memory there is.

You can also see the size of the L2 cache installed in your system by using the `lsattr` command:

```
lsattr -El L2cache0
```

3.5.3 PCI Adapter Configuration

If the required device drivers are installed, PCI adapters are configured automatically at system restart or by using the `cfgmgr` command.

When you install the AIX operating system, it will query the PCI bus to identify all AIX supported adapters. The installation process will install the required device drivers and diagnostic support (if available), and the system will automatically configure the PCI adapters.

If you have a system installed and running, you can use `smit cfgmgr`. This method will require you to put the AIX installation media in place (CD-ROM, tape, and so forth) before configuring the adapters. The media will be used only if the device drivers for your adapters are not currently installed.

You can also manually install the device drives and diagnostics before running `cfgmgr` or rebooting the system for automatic PCI adapter configuration.

When using PCI adapters, the slot number is important. If you change an adapter from slot A to slot B, the original adapter information for slot A will be kept in a `defined` state; a new definition will be created for slot B, and a new device will be configured.

3.5.3.1 Configuring Non-Graphic PCI Adapters

Configuring PCI adapters is very easy because they are actually configured automatically when their driver software is installed and the `cfgmgr` command is run.

The easiest way to configure non-graphic PCI adapters (such as LAN adapters, and so on) installed after IPL is to proceed using SMIT, enter:

```
smit devices
-> Install/Configure Devices Added After IPL
```

You may have to provide a source for the device driver software package for the new PCI device if it is not yet installed.

Note

Do not use the above procedure if you want to add or replace graphic adapters because if anything goes wrong (for example, a required device driver is not yet installed, and so forth), you will disable your graphics console.

To add or change PCI graphic adapters, proceed as described in 3.3, “Graphics Accelerators” on page 38.

If, for whatever reason, you need to install software but you are unable to boot from hard disk, or you have no working display, you may install the software while in maintenance mode.

3.5.3.2 Configuring SCSI Devices

The RS/6000 Model B50 has one integrated Ultra SCSI Interface.

The process of configuring SCSI devices on the 7046 Model B50 is exactly the same as it is on other RS/6000 Systems. This means you can use the SMIT support to configure devices that are added after IPL enter the `cfgmgr` command, or reboot the system. However, remember that unlike with AIX Version 3.2, on AIX Version 4, the software support for the device you are adding may not yet be installed. Therefore, check to see if the device driver for your SCSI device is installed. If not, use SMIT to configure the device and provide an input device or directory where the installation software can be accessed.

The Web-Based System Management tool, wsm, provides a similar function as SMIT, but using a visually rich GUI and task-oriented dialogs.

3.6 Summary Chart of Adapters, Size, and Placement

Table 12 summarizes the adapters (communication, storage, and graphics) available on the RS/6000 7046 Model B50 system and describes the adapter placement guidelines.

Table 12. Communication, Graphics, and Storage Adapters Summary

Adapter	Max per System	Slot Allocation	Feature Code	Size
4/16 High-Speed Token Ring PCI	2	1-2	# 4959	short
10/100 Mbps Ethernet PCI	2	1-2	# 2968	short
ISDN Basic Rate PCI	1	1-2	# 2708	short
SysKonnect SK-NET FDDI-LP SAS PCI	1	1-2	# 2741	short
SysKonnect SK-NET FDDI-LP DAS PCI	1	1-2	# 2742	short
SysKonnect SK-NET FDDI-UP SAS PCI	1	1-2	# 2743	short
TURBOWAYS 155 UTP ATM PCI	2	1-2	# 2963	short
TURBOWAYS 155 MMF ATM PCI	2	1-2	# 2988	short
8-Port Asynchronous EIA-232/RS-422 PCI	2	1-2	# 2943	short
128-Port Asynchronous Controller PCI	2	1-2	# 2944	short
2-Port Multiprotocol PCI	2	1-2	# 2962	short
ARTIC960Hx 4-Port Selectable PCI	1	2	# 2947	long
POWER GXT130P	1	1	# 2830	short

Adapter	Max per System	Slot Allocation	Feature Code	Size
PCI Single-Ended Ultra SCSI	2	1-2	# 6206	short
PCI Differential Ultra SCSI	2	1-2	# 6207	short
PCI Dual-Channel Ultra2 SCSI Adapter	1	1-2	# 6205	short
PCI 3-Channel Ultra2 SCSI RAID	1	2	# 2494	long
PCI Advanced SerialRAID	1	2	# 6225	long

Chapter 4. Storage: Architectures and Devices

In this chapter, the internal storage features of the Model B50 are described. Storage standards, SCSI, and SSA are covered. The different technologies are discussed, and additional information is provided in the cabling sections. In addition, a description of the storage devices and adapters follows.

4.1 Storage Boot Devices

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

- SCSI CD-ROM
- SCSI DASD devices connected to a base SCSI controller
- SCSI tape drives
- SSA external DASD (if the machine was installed using NIM)

4.2 SCSI Overview

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

Note

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

4.2.1 SCSI-I

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

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

4.2.2 SCSI-II

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

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

4.2.2.1 Common Command Set

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

4.2.2.2 Tagged Command Queuing

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

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

4.2.2.3 Disconnect/Reconnect

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

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

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

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

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

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

4.2.3 SCSI-III

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

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

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

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

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

4.2.3.1 Overview of SCSI-III Standards

The breakdown of the SCSI-III standards is shown in Figure 10 and is described in Table 13.

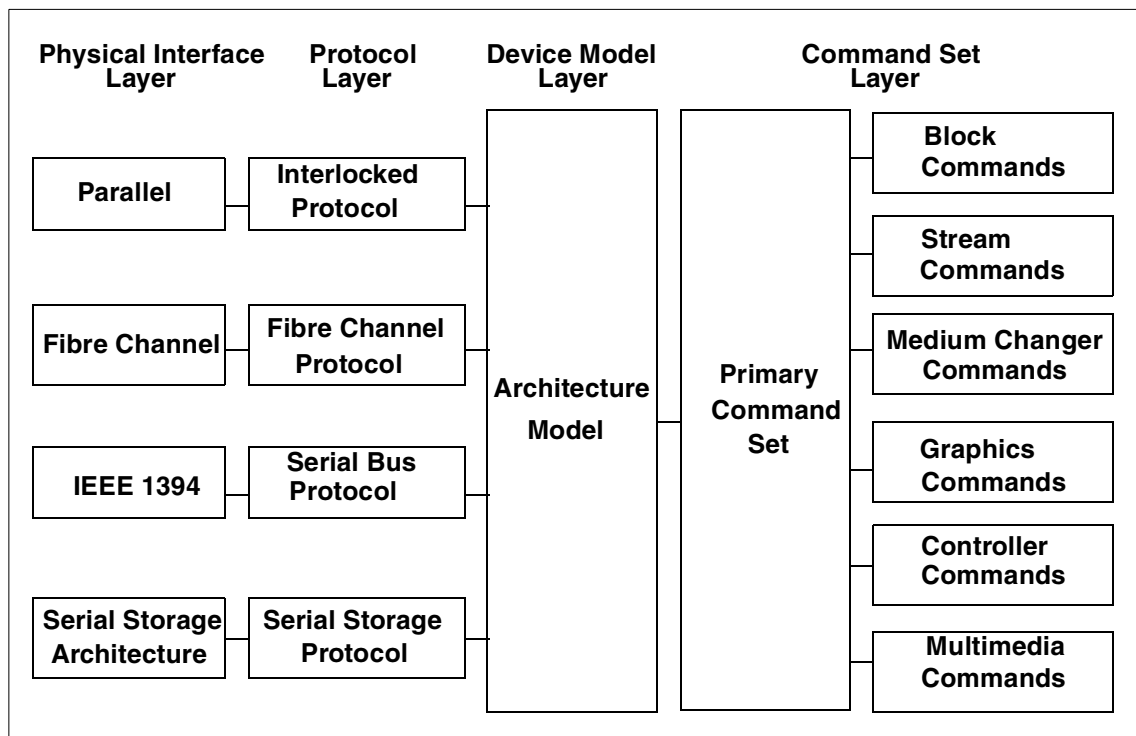


Figure 10. SCSI III Standards Overview

Table 13. Overview of SCSI-III Standards

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

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

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

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

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

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

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

Note

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

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

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

4.2.4 SCSI Terminology

The terminology of SCSI products can sometimes be confusing. Listed in the following sections are some definitions to help your understanding.

4.2.4.1 General Terminology

The following is a list of general terminology for SCSI:

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

Wide	Wide is used in reference to the width of the SCSI parallel bus between the adapter and the device. Wide means wider than the original 8-bit path defined in SCSI-I, usually 16-bit. 32-bit transmission is possible within the specification but there are no Wide-32 devices on the market. With a 16-bit path, the data rate is double that of an 8-bit device for the same clock speed.
Fast/Wide	Fast/Wide refers to a 16-bit data path running at 10 MHz producing a maximum data transfer rate (or <i>burst</i> rate) of 20 MB/s.
Fast-20	Fast-20 is a bus running at double the clock speed of Fast, or 20 MHz. Fast-20 typically refers to an 8-bit bus and can also be called <i>Fast-20 Narrow</i> . <i>Fast-20 Wide</i> is the 16-bit version, also known as Ultra SCSI.
Ultra SCSI	Ultra SCSI, as described above, is a subset of the SCSI-III specification. It is effectively a Fast SCSI bus running at 20 MHz. Ultra SCSI can produce a maximum transfer of 20 MB/s over an 8-bit data path. Wide Ultra SCSI, the 16-bit version of Ultra SCSI can transmit a maximum of 40 MB/s.
Ultra2 SCSI	Ultra-2 is the latest extension to the SCSI-III specification that allows transfers to take place at 80 MB/s on a wide bus, or 40 MB/s for narrow. At twice the speed of the current Ultra SCSI, and comparable with SSA, it has the advantage of complete backwards compatibility with existing peripherals. The cable length maximum, as defined by the standard, is 12 meters. This standard requires Low Voltage Differential (LVD).

Note

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

4.2.4.2 Asynchronous versus Synchronous

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

and to the devices. This process prevents data from being lost and ensures that data transmission is error free.

4.2.4.3 SCSI Differential

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

Consider Figure 11 as an example. Signal A is transmitted along two wires, A+ and A-. A- is the mirror image of A+. When some line noise appears, such as a signal spike, the spike will be of the same polarity on both A+ and A-. On the receiving end, the signals are subtracted, and the result is A+/A- as shown in the diagram. The spike is canceled out by the subtraction.




Figure 11. SCSI Differential Example

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

4.2.5 SCSI Repeaters

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

As noted in Table 14, the maximum length of the SCSI bus for a SCSI-II F/W configuration is 3.0 meters. With a repeater, it is possible to extend this distance.

4.2.6 Summary of SCSI Specifications

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

Table 14. Differential SCSI Modes

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

Note

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

4.3 SCSI Cabling

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

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

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

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

The following describes the SCSI cabling of the two PCI SCSI adapters supported by the Model B50 entry rack mount server:

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

Note

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

4.3.1 General SCSI Cabling Considerations

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

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

4.3.1.1 Default SCSI Cable Routing

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

4.3.1.2 SCSI Bus Length

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

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

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

Devices that have two SCSI connectors have internal cabling that must be included when calculating the total cable length. When attaching these devices to a SCSI chain, connect one cable to one connector and the other cable to the other connector. Do not piggy back the second cable/terminator onto the first as you would with a device that had only one SCSI connector as shown in Figure 12.




Figure 12. Piggy Back Connectors

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

4.3.1.3 SCSI Terminators

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

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

external bus and that the end of the internal SCSI cable has been properly terminated.

4.3.1.4 SCSI Device Addresses

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

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

Note

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

- If a SCSI address is changed after the operating system is loaded, the operating system must be stopped and loaded again to have the correct configuration, or `cfgmgr` run.
- Standalone Diagnostics always default to a SCSI ID of 7 when testing SCSI adapters and devices. Choosing SCSI IDs other than 7 for both adapters prevents problems when using Standalone Diagnostics on systems in HA clusters or in multi-initiator configurations.

- Check the documentation for your specific SCSI subsystem to insure that there are no SCSI ID conflicts if the adapters are addressed at ID(s) other than 6 or 7.

4.3.1.5 SCSI Bus Width

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

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

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

4.3.2 Cabling the PCI Single-Ended Ultra SCSI Adapter

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

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

For this adapter:

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

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

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

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

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

Cables and Terminators for Single-Ended SCSI Adapters

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

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

Adapter-to-First Device Cables

Table 15 describes the cables for the Single-Ended Ultra SCSI Adapter for connection to the first device.

Table 15. Single-Ended Ultra SCSI Adapter-to-First Device Cables

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

Note

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

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

Device-to-Device Cables

Table 16 describes the device-to-device cables for single-ended applications.

Table 16. Device-to-Device Cables for Single-Ended Installations

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

Note

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

System-to-System Cables

Table 17 on page 75 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 17. 16-bit SCSI-II System-to-System Cable

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

Terminators for Use with These Adapters

Table 18 describes the terminators for single-ended installations.

Table 18. Terminators for Single-Ended Installations

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

4.3.3 Cabling the PCI Differential-Ended Ultra SCSI Adapter

To understand the cabling for the PCI differential-ended Ultra SCSI adapter (# 6207), first read 4.3.1, "General SCSI Cabling Considerations" on page 68 and then read the following for specific information.

SCSI Differential Cable Lengths Using This Adapter

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

Adapter-to-First Device Cables

Table 19 provides a list of adapter to first device cables.

Table 19. Differential-Ended Ultra SCSI Adapter-to-First Device Cables

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

Note

When cables are ordered by feature code, the appropriate terminator is included with the order. When cables are ordered by part number, only the cable is included. For terminator part numbers, refer to Table 21 on page 78.

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 20 provides a list of device-to-device cables.

Table 20. Device-to-Device Cables for Differential-Ended Installations

M/T	Feature	Part Number	Length (Meters)	Cable Description
DE External Device	2848 /9134	74G8511	0.6	Device-to-device (where second device has two connectors), 8-bit narrow bus

M/T	Feature	Part Number	Length (Meters)	Cable Description
DE External Device	2860 /9139	52G9921	0.3	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2884 /9160	52G4291	0.6	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2846 /9132	52G4233	2.5	Device-to-device (where second device has two connectors), 16-bit wide bus
7027	2425			
DE External Device	2885 /9161	88G5749	4.5	Device-to-device (where second device has two connectors), 16-bit wide bus
7027-HSD	3132	40H7351	6.0	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2870 /9146	88G5747	12.0	Device-to-device (where second device has two connectors), 16-bit wide bus
7027 HSD	3125			
DE External Device	2869 /9145	88G5748	14.0	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2868 /9144	88G5746	18.0	Device-to-device (where second device has two connectors), 16-bit wide bus
7027 HSD	3136			

Note

When cables are ordered by feature code, the appropriate terminator is included with the order. When cables are ordered by part number, only the cable is included. For terminator part numbers, refer to Table 21 on page 78.

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

Terminators for Use with This Adapter

This adapter has on-card SCSI terminators that must be removed before the adapter can be used in a high-availability configuration. The high-availability configuration is implemented by removing the three on-card differential terminating resistors (labeled RN1, RN2, and RN3) on the adapter, then attaching the middle leg connector of the high-availability configuration Y-cable to the adapters 68-pin external connector. The remaining two legs of the Y-cable are used to attach other systems and devices to the SCSI bus.

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

Note

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

Table 21 provides a list of required terminators.

Table 21. Terminator for Description-Ended Installations

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

Cabling Examples for the PCI Differential Ultra SCSI Adapter

Figure 13 on page 79 and Figure 14 on page 79 show how to cable the PCI Differential Ultra SCSI Adapter.

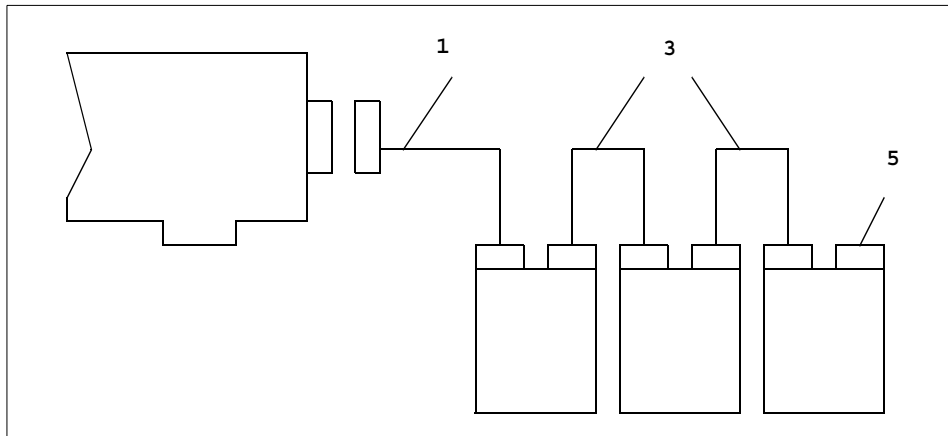


Figure 13. Differential External Narrow Bus




Figure 14. Differential External Wide Bus

Table 22 provides a list of available cables for differential adapters.

Table 22. Cabling for the PCI Differential Ultra SCSI Adapter

Item Number	Part Number	Length (Meters)	Description
1	06H6037	1.0	Adapter-to-dual-connector device (narrow 8-bit)

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

4.3.4 Cabling the PCI Dual-Channel Ultra2 SCSI Adapter (# 6205)

To understand the cabling for this adapter, read 4.3.1, "General SCSI Cabling Considerations" on page 68, then read the following for specific information.

Supported Ultra2 SCSI Configurations for Each Channel

Rules for external devices running at Ultra2 speeds in a 2104 Expandable Storage Plus subsystem are as follows:

- One dual-port 2104 Expandable Storage Plus subsystem can be attached. Maximum cable length must not exceed 25 meters (approximately 80 feet).
- One dual-port 2104 Expandable Storage Plus subsystem can be attached with its second port driven from another PCI Dual-Channel Ultra2 SCSI Adapter so that the drives in the subsystem are shared between the two systems.

Terminators for Low-Voltage Differential Mode of This Adapter

The 2104 Expandable Storage Plus subsystem contains Low-Voltage Differential (LVD) terminators that operate in single-ended mode if any single-ended devices are on the SCSI bus. Otherwise, the bus operates in LVD mode at Ultra2 speeds (80 MB/sec).

4.4 Serial Storage Architecture Overview

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

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

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

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

4.5 Comparison between SCSI and SSA Architecture

Table 23 lists the common features found in SCSI and SSA devices for comparison and Figure 15 shows how the subsystems differ.

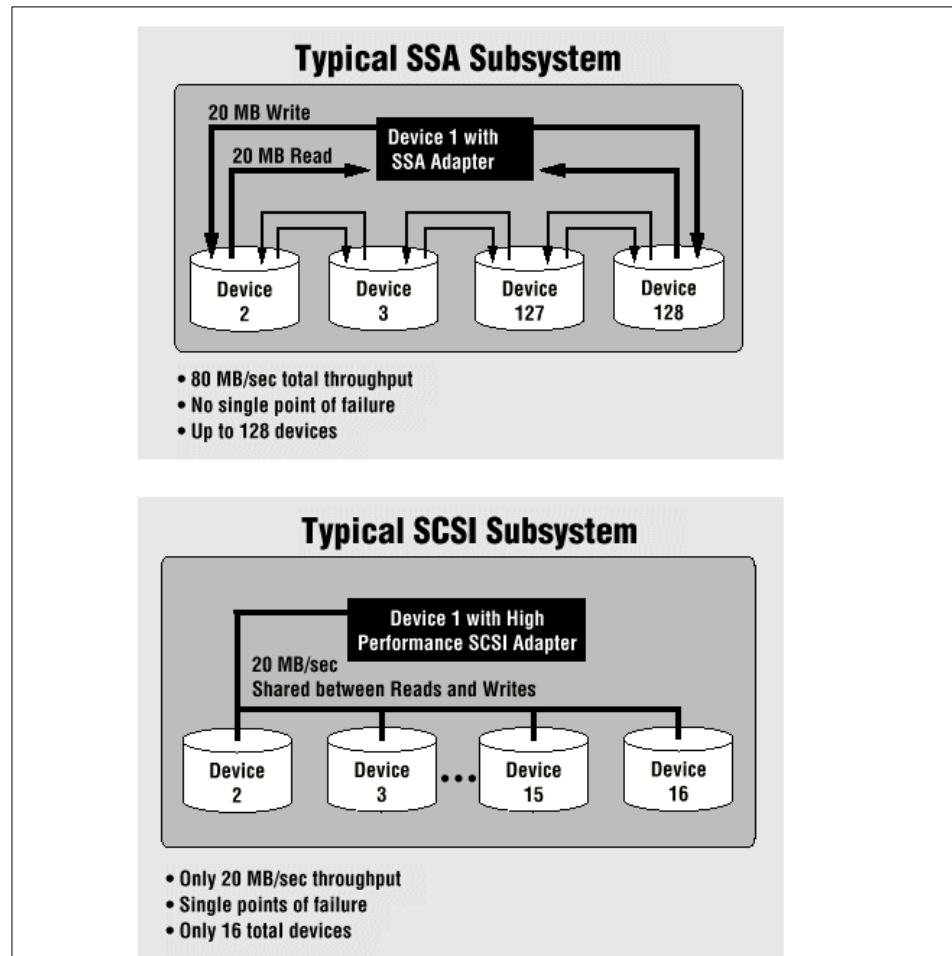


Figure 15. Comparison between SCSI and SSA Subsystem

Table 23. Comparing SSA and SCSI

Feature	SCSI	SSA
Cable Type	68 conductor shielded cable (SCSI-II Differential Fast/Wide)	Four conductor twisted-pairs

Feature	SCSI	SSA
Maximum Cable Lengths	25 meters total	25 meters between devices, up to 2500 meters between devices with the use of a Fiber Optic Extender
Cable Costs	Significant	Low
Individual Drive Addressing	Must be set manually	Dynamic
Cable Termination	Electrical terminator must be at the end of the SCSI-bus	No termination required, closed serial loops
Fault Tolerance	If single data path (bus) is compromised, access is lost	No single point of failure. Multiple data paths prevents a single loop from causing loss of access to information
Maximum Throughput	Up to 40 MB/s	Up to 160 MB/s

4.6 Internal SCSI Devices

The following section gives an overview of the internal SCSI devices for the Models B50 and their main characteristics:

- Disk drives
- CD-ROM drives
- Tape drives

4.6.1 Disk Drives for the Model B50

Table 24 on page 84 shows the main characteristics of the internal disk drives that are available for the Model B50. You can install combination of a 9.1 GB (# 3112) or a 18.2 GB (# 3113) hard disk up to two drives into internal disk drive bays.

4.6.1.1 9.1 GB Ultra SCSI 16-Bit Disk Drive (# 3112)

The 9.1 GB Ultra SCSI 16-bit 1-inch (25 mm) disk drive is a high performance disk drive that provides 9.1 GB capacity that supports the Ultra SCSI interface speed of 40 MB/s.

4.6.1.2 18.2 GB Ultra SCSI 16-bit Enhanced Disk Drive (# 3113)

The 18.2 GB Ultra SCSI 16-bit 1-inch (25 mm) High Disk Drive offers high-performance that provides 18.2 GB of storage capacity and supports the industry Ultra SCSI interface speed of 40 MB/s.

Table 24. Internal Disk Drives for the Model B50

Capacity	18.2 GB	9.1 GB
Feature Code	# 3113	# 3112
Maximum Media Data Transfer Rate	244.0 Mb/s	171.1 Mb/s
Average Seek Time	7 ms	7.5 ms
Rotational Speed	7200 RPM	7200 RPM
Average Latency	4.17 ms	4.17 ms
SCSI Interface	Ultra SCSI	Ultra SCSI
Data Width	16-bit, single-ended	16-bit, single-ended

4.6.2 Internal CD-ROM Drive

Table 25 provides an overview of the supported CD-ROM drive for the Model B50. The 32X max SCSI-II Internal CD-ROM drive is a tray loading CD-ROM drive providing up to 4800 KB/sec maximum media data transfer rate. It is a 5.25 inch, half-high form factor, single-ended, 8-bit, multi-session capable, CD-ROM drive that provides state-of-the-art performance and supports existing 600 MB CD-ROM discs. It will operate in either the vertical or horizontal positions.

Table 25. Internal CD-ROM Drives for Model B50

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

4.7 Internal SSA Devices

There are no internal SSA storage devices available for the Model B50. However, you can attach external SSA devices through the PCI Advanced SerialRAID Adapter (# 6225).

For more information about this adapter, see 3.2.2, “SSA Adapter” on page 37.

4.8 External SCSI and SSA Storage Devices

For the Model B50, two primary rack-mounted storage solutions are available offering SCSI and SSA architectures. This section describes the Expandable Storage Plus and the IBM Model 7133 Serial Disk System. There are also many external SCSI and SSA storage devices available that can be attached to the Model B50. See also 2.1.3.11, “Supported Devices” on page 23 for details about the Model B50 supported devices.

Further information can be found on the Web at:

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

4.8.1 Expandable Storage Plus (# 2104)

Figure 16 shows the IBM Expandable Storage Plus Model DL1 and Model TL1.



Figure 16. IBM Expandable Storage Plus SCSI

The IBM Expandable Storage Plus disk enclosure provides low cost disk storage for RS/6000 servers. In addition, Expandable Storage Plus offers space for up to ten hot swap Ultra 2 (LVD) disk drives, including IBM's high performance 9 and 18 GB Ultrastar drives (both 7,200 rpm or 10,020 rpm disk drives are available) and 36 GB drive (7,200 rpm disk is available).

It is designed to meet the needs of RS/6000 customers that have external disk capacity requirements and desire to use a Ultra 2 (LVD) interface. When combined with the IBM PCI 3-channel Ultra SCSI RAID Adapter, high performance and high availability are available. IBM's Expandable Storage Plus provides leading edge technology at a very affordable price.

IBM intends to provide HACMP support for the IBM Expandable Storage Plus when used with RS/6000 servers and the appropriate adapters.

4.8.1.1 Supporting Adapters

The IBM Expandable Storage Plus attaches to RS/6000 servers with the following adapters:

- IBM SCSI-2 Fast/Wide Adapter/A (# 2415)
- IBM PCI RAID Adapter (# 2493)
- PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494)
- PCI Dual Channel Ultra2 SCSI Adapter (# 6205)
- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)

The appropriate cable to connect the adapter and Expandable Storage Plus is also required.

4.8.1.2 At a Glance

The IBM Expandable Storage Plus provides:

- Disk Enclosure for up to 10 high-performance hot swap IBM Ultra 2 (LVD) disk drives, including both 7,200 and 10,020 rpm options
- Capacity ranging from 9 GB to 360 GB, utilizing 9, 18, or 36 GB disk drives
- Redundant power and cooling
- Advanced monitoring using SCSI Enclosure Services
- Rack and Tower models
- Support across a wide variety of RS/6000 servers
- IBM SystemXtra support services and financing

4.8.2 IBM 7133 Serial Disk System Advanced Models T40 and D40

Figure 17 shows the IBM 7133 Serial Disk System Advanced Models T40 (deskside tower) and D40 (rack-mountable drawer).



Figure 17. IBM 7133 Model T40 Deskside Tower and 7133 Model D40 Drawer

For the latest in high-throughput and high-density IBM storage for UNIX-based servers, you need the 7133 Serial Disk System models T40 and D40. Using IBM Serial Storage technology, up to 96 36.4 GB disk drives can be connected using a single SSA Adapter.

Each 7133 can be shared by up to eight servers, enabling you to share disks and create a very high-performance, high-availability environment for running your business applications.

- One of the highest performing IBM disk systems.
- No single point of data-path failure.
- Up to sixteen 9.1, 18.2, or 36.4 GB hot-swappable SSA disk drive modules.
- Concurrent maintenance of redundant power and cooling.
- Connects up to eight servers.
- Up to 10 km between disk systems with the Advanced SSA Optical Extender.
- Up to 35 MB/sec sustained data throughput per SSA adapter.
- Compliant with the SSA interface specifications (ANSI X3T10.1).

- Attaches to supported models of Sun servers using the IBM SSA Interface Controller Card for Sun SBus (# 4003).
- Attaches to supported models of HP, Sun, and DEC servers using the 7190 Ultra SCSI Host to SSA Loop Attachment.
- Attaches to selected Intel-based servers using the IBM SerialRAID Adapter.
- Models D40 and T40 may be intermixed in an SSA loop with other 7133 models and the 7131-405.

Table 26 summarizes the key features of the 7133 Serial Disk System Advanced Models D40 and T40.

Table 26. IBM 7133 Serial Disk System Advanced Models at a Glance

Disk Storage Capacity	(Min-max GB per disk system or adapter)			
	Disk drive capacity	9.1 GB	18.2 GB	36.4 GB
	Per disk system	36.4-145.6	72.8-291.2	145.6-582.4
	Per adapter	36.4-873.6	72.8-1752.0	145.6-3494.4
System Capacity	Up to six disk systems per adapter Up to 96 drives per adapter			
Currently Available RS/6000 Adapters	Advanced SerialRAID Adapter (# 6225)			
Operating Systems	AIX Version 4 The specific AIX release level requirements are RS/6000 host system and SSA adapter dependent.			
Other Supported UNIX-based Servers	Sun SPARCstation 20; SPARCserver 1000 and 1000E SPARCcenter 2000 and 2000E Ultra Enterprise 2, 3000, 4000, 5000, and 6000 HP 9000 Series 800 HP 9000 Enterprise servers including D, K, T, G, H, and I class systems, Digital AlphaServer 300, 800, 1000, 1000A, 2000, 2100, 2100A, 4000, 4100, 8200, and 8400 Series			
Supported Intel-based Servers	IBM PC Server models 325 and 330 IBM Netfinity 5000, 5500, 5500 M10, 7000, and 7000 M10 Compaq ProLiant 1500, 2500, 5000, 5500, 6000, 6500, and 7000, HP NetServer LH Pro, LX Pro, LHx, LXr, and Pro 8			

Multiple Systems Attachment	Up to eight with F/C #6225 or #6230 (non-RAID) Up to two with F/C #6235 (RAID 5) Up to four with 7133 F/C #4003 Up to four with 7190 Ultra SCSI Host to SSA Loop Attachment				
Warranty Period	One-year limited warranty				
Dimensions	<table> <tr> <td>Model D40</td> <td>Rack mountable drawer, 4EIA units 6.7"H x 17.5"W x 29"D (171 mm x 444 mm x 737 mm)</td> </tr> <tr> <td>Model T40</td> <td>Deskside Tower 24.0"H x 8.3"W x 33.9"D (610 mm x 210 mm x 861 mm)</td> </tr> </table>	Model D40	Rack mountable drawer, 4EIA units 6.7"H x 17.5"W x 29"D (171 mm x 444 mm x 737 mm)	Model T40	Deskside Tower 24.0"H x 8.3"W x 33.9"D (610 mm x 210 mm x 861 mm)
Model D40	Rack mountable drawer, 4EIA units 6.7"H x 17.5"W x 29"D (171 mm x 444 mm x 737 mm)				
Model T40	Deskside Tower 24.0"H x 8.3"W x 33.9"D (610 mm x 210 mm x 861 mm)				
Ordering Highlights Base unit does not include disk drives. Add any of the supported disk drive modules for a total of four to sixteen disk drives. Two power supplies are required.					
Feature Code	Description				
8031	Raven-black cover for model D40				
8032	Pearl-white cover for model D40				
8209	One 9.1 GB advanced disk drive module				
8218	One 18.2 GB advanced disk drive module				
8236	One 36.2 GB advanced disk drive module				
8801	One 1 meter advanced SSA cable				
8802	One 2.5 meter advanced SSA cable				
8805	One 5 meter advanced SSA cable				
8810	One 10 meter advanced SSA cable				
8825	One 25 meter advanced SSA cable				
8851	Advanced SSA Optical Extender (pair)				
8852	Advanced MultiMode 50 Extender (pair)				
8853	Advanced MultiMode 62 Extender (pair)				

Chapter 5. Hardware Architecture Overview

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

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

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

5.1 Hardware Design for the RS/6000 7046 Model B50

The RS/6000 Model B50 is a uniprocessor system that utilizes the PowerPC 604e processor running at 375 MHz. It supports an integrated 1 MB parity-checked synchronous L2 cache. The Model B50 also features an enhanced memory controller that uses ECC SDRAM memory.

The local system bus on the RS/6000 Model B50 is a 60X bus running at 83 MHz. This bus is 64 bits wide (with an additional 8 bits for parity) and is synchronous. The processor clock and regulator card reside as one unit, namely the OVdd/Vdd/PLL card. This is illustrated in the planar design of the Model B50 in Figure 18 on page 92.

In the Model B50, the 60X bus runs from the PowerPC 604e processor and connects to an enhanced memory controller chip that 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 B50, 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.

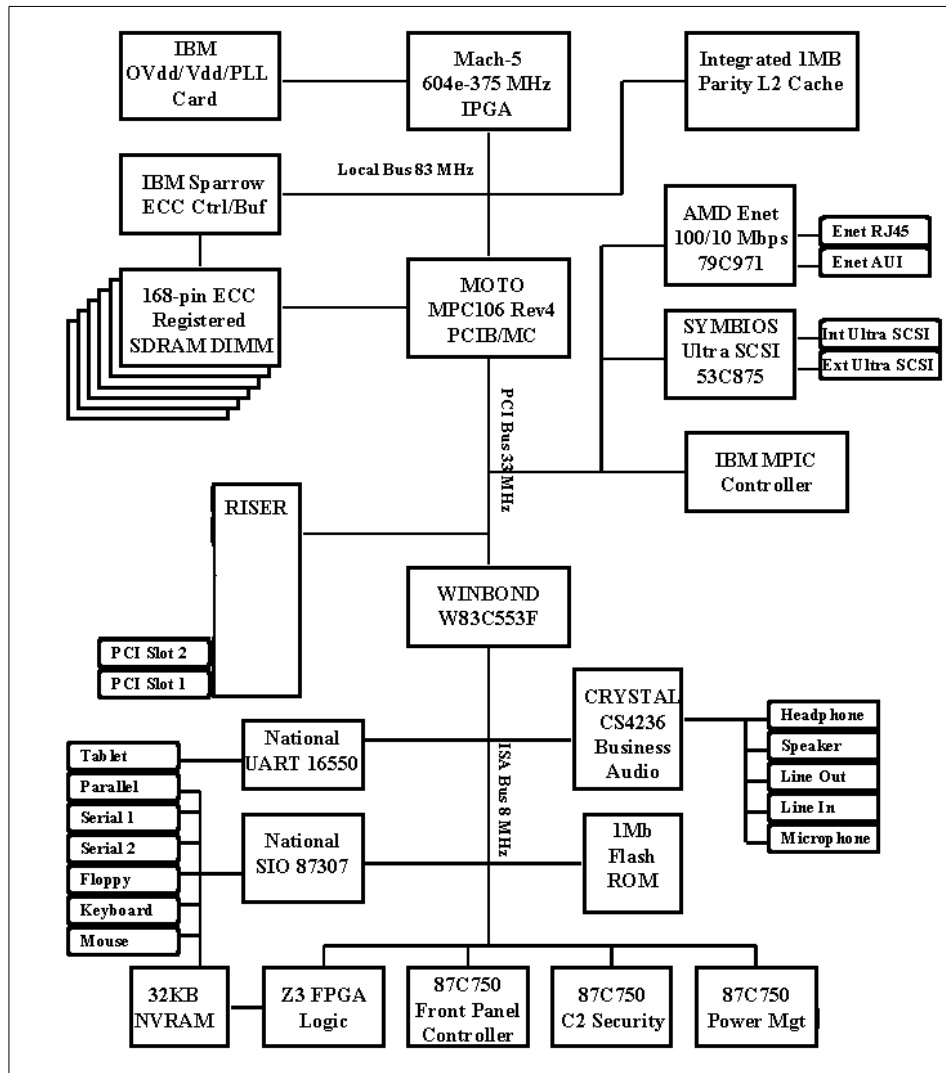


Figure 18. Planar Block Diagram of the RS/6000 Model B50

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 that the Model B50 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)

5.2 Bus Architecture

The following sections cover the different bus architectures that exist on the RS/6000 7046 Model B50.

5.2.1 The PCI Bus Architecture

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

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

5.2.1.1 PCI Specification Revisions

According to PCI Specification Revision 2.0, the PCI bus operates on 32 or 64-bits of data at a clock speed of 33 MHz. However, as the speed of processors continues to increase and devices such as video adapters require more bandwidth, bus speeds must also increase to avoid data bottlenecks. Consequently, PCI Specification Revision 2.0 has been superseded by PCI Local Bus Specification Revision 2.1, which introduces support for a 66 MHz bus operation speed, doubling the throughput of the former 33 MHz standard. The RS/6000 Model B50 fully complies with PCI Local Bus Specification Revision 2.1 but only supports 33 MHz 32-bit slots.

Under PCI Specification Revision 2.1, the 66 MHz bus operation speed is an optional feature and is backward-compatible with 33 MHz PCI devices and

buses. If a 66 MHz capable PCI device is installed in a 33 MHz PCI bus, the device must operate at 33 MHz. Likewise, if any 33 MHz PCI devices are installed into a 66 MHz PCI bus, the PCI bus must operate at 33 MHz.

5.2.1.2 PCI Features and Benefits

The PCI bus architecture has many advantages including the following:

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

5.2.2 The ISA Bus Architecture

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

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

On the RS/6000 7046 Model B50 computer, the ISA bus is used to attach internal devices such as the diskette drive. No ISA slots are provided for the addition of non-native ISA devices.

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

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

Table 27 provides the bus specification for different architectures and compares them to the PowerPC processor's speed.

Table 27. PowerPC and Bus Specification

Specification	PowerPC	PCI	ISA(8)	ISA(16)	EISA
Processor Speed	166, 200 233, 375				
Data Bus	64	64	8	16	16/32
Address Bus	32	64	20	24	24/43
Bus Clock	66 MHz 83 MHz 100 MHz	33 MHz (32-bit)	4.7 MHz	8.3 MHz	8.3 MHz
Interrupts		4	6	11	11
DMA Channel		busmaster	3	7	7

5.2.3 The System Buses

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

The 60X bus uses separate address and data buses and a variety of control and status signals. The address bus is 32 bits wide (with 4 bits parity) and the data bus is 64 bits wide (with 8 bits parity). These separate address and data buses run independently, allowing split transactions. The bus is synchronous with all timing relative to the rising edge of the bus clock. Inputs are sampled, and outputs are driven from this edge. It is through this bus that the two-way SMP processors are attached, memory is accessed, and cache snooping is achieved.

5.2.4 The 6XX-MX I/O Bus

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

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

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

5.2.5 The X-Bus

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

- The Keyboard/Mouse Controller
- Real-Time Clock and Non-Volatile RAM

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

- Mini-Support Processor (I²C Controller)

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

5.3 The Processor Subsystem

Since its inception, the RS/6000 product line has used a number of different processor designs, the original implementation being the POWER architecture.

5.3.1 The PowerPC 604e Microprocessor

The RS/6000 7046 Model B50 features the PowerPC 604e microprocessor. The 604e is an equivalent, but enhanced, derivative of the 604. Enhancements to the PowerPC 604e, exceeding its predecessor (PowerPC 604), include:

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

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

Figure 19 shows the PowerPC 604e microprocessor architecture.

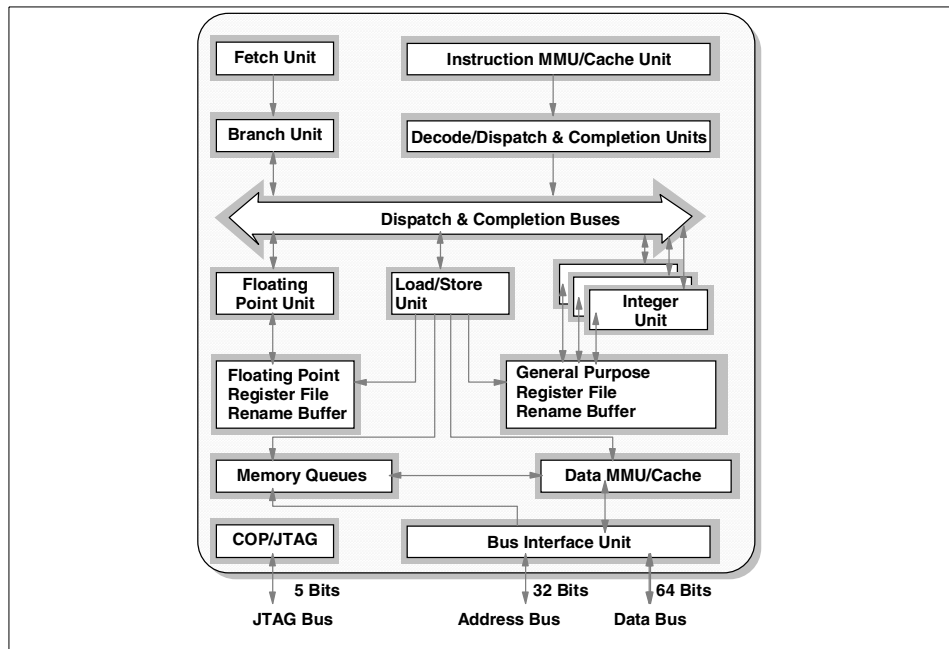


Figure 19. PowerPC 604e Microprocessor Logical Block Diagram

A 604e is defined by the following specifications:

- PowerPC 604e microprocessor running at:
 - 375 MHz on the RS/6000 Model B50
- Superscalar design with integrated integer, floating-point, and branch units
- 32 KB four-way set-associative instruction cache
- 32 KB four-way set-associative data cache
- 64-bit memory interface with 32-bit addressing
- Virtual memory support for up to 4 petabytes (2^{52})
- Real memory support for up to 4 gigabytes
- ~5.1 million transistors

- Nap power management mode
- JTAG bus (IEEE 1149.1) for RAS

5.4 The RS/6000 Platform Architecture

The RS/6000 Platform Architecture (RPA) introduces a comprehensive computer system hardware-to-software interface definition, combined with minimum system requirements, that enables the development and porting of software to a range of compatible computer systems from portables through servers. These systems are based on the PowerPC microprocessor, as defined in the *PowerPC Architecture: A Specification for a New Family of RISC Processors*, ISBN 1-5586-0316-6. The RS/6000 7046 Model B50 is RPA compliant.

5.4.1 The RS/6000 Platform Architecture Introduction

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

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

A key benefit of the RPA specification is the ability of hardware platform developers to have degrees of freedom of implementation below the level of architected interfaces and, therefore, have the opportunity for adding unique value. This flexibility is achieved through architecture facilities including:

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

5.4.2 RS/6000 Platform Architecture Achievements

The RS/6000 Platform Architecture specification achieves the following:

- Creates a stable platform architecture to be used by all RS/6000 platforms based on PowerPC processors and PCI I/O bus (and may contain other I/O buses in addition to the PCI bus).
- Creates an architecture that allows platforms to operate with a previous version of AIX (n-1 capability).
- Leverages existing and future industry-standard buses and interfaces. Existing bus architectures have a proven level of performance and function. Established industry-standard interfaces (for example, SCSI, IDE, LocalTalk, and Ethernet) and newer bus architectures, interfaces, and protocols (PCI, PC Card, IrDA, to name a few) provide higher levels of performance or utility that are not achievable by the older standards. The architecture allows platform and system designers to determine which buses, interfaces, and protocols best suit their target environment.
- Provides a flexible address map. Another key attribute of this specification is the relocatability of devices and subsystems within the PowerPC address space. Subsystem address information, which defines where I/O devices reside, is detected by the Open Firmware and passed to the operating systems in the device tree. The architecture accommodates the use of multiple identical buses and adapters in the same platform without address conflicts.
- Builds upon the Open Firmware boot environment defined in *IEEE 1275, IEEE Standards for Boot (Initialization Configuration) Firmware, Core Requirements and Practices*, IEEE part number DS02683, ISBN 1-5593-7426-8. To order copies of IEEE standards, load the following Web site: <http://standards.ieee.org/index.html>.

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

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

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

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

Chapter 6. Diagnostics and System Management Services

One of the most important features provided by the PCI-based RS/6000 workstations and servers is support of the AIX diagnostics programs. Workstations, and especially servers, should, by their nature, provide enhanced Reliability, Availability, and Serviceability (RAS) features. Access to AIX diagnostics, Standalone and Online, helps provide these features.

In addition to providing diagnostic support for new PCI devices available for the PCI-based RS/6000 systems, a major change had to be made to the firmware. An extra function had to be added to the firmware to allow a logical representation of the keyswitch that would allow the system to either boot in normal or service mode. The Model B50 has an electronic keyswitch function and the system can either be booted in normal or service mode by pressing a function key during the firmware POST stage. Additionally, support was added to enable the system to load diagnostics from a CD-ROM.

Two methods of running diagnostics are supported on the PCI-based RS/6000 systems:

- Standalone Diagnostics
- Online Diagnostics

Both of these methods are described in this chapter.

6.1 Diagnostics Operating Considerations

The following items identify some issues to consider before using the diagnostics programs:

- **Device Diagnostics**

If diagnostics support is not installed on your system, and you install this option, device support for devices physically installed in your machine will not be automatically installed. Diagnostics support is, however, automatically installed when you install from an AIX Version 4 CD-ROM (check with the `ls1pp -L` command if `bos.diag.*` and `bos.rte.diag` filesets are already installed on your system since they are required). The Online Diagnostics test list will display entries for all the devices in your system that have diagnostic support.

- **Selecting a Console Display**

When you run Standalone Diagnostics, you will have to select the console display. The diagnostic program displays instructions on graphics displays and on terminals attached to the S1 serial port.

- **Identifying the Terminal Type to Diagnostics**

When you run diagnostics, the diagnostic program must know what type of terminal you are using. If the terminal type is not known when the Function Selection menu is displayed, the diagnostics will not allow you to continue until a terminal is selected from the Define Terminal option menu.

- **Undefined Terminal Type**

If an undefined terminal type from the Define Terminal option menu is entered, the menu will prompt you to enter a valid terminal type, and the menu will be redisplayed until a valid type is entered or you exit from the option.

- **Resetting the Terminal**

If you enter a valid terminal (according to the Define Terminal option), but it is not the correct type for the ASCII terminal being used, difficulty may be encountered in reading the screen, using the functions keys, or the Enter key. These difficulties can be bypassed by pressing **Ctl-c** to reset the terminal. The screen display that results from this resetting action varies with the mode in which the system is being run:

- Online Concurrent or Single-User Mode - The command prompt appears.
- Standalone Mode or Online Service Mode - The terminal type will be reset to dumb, the diagnostics operation instruction panel will be displayed, and you will be required to go through the Define Terminal process again.

- **Running Diagnostics from a TTY Terminal**

The following should be considered when running diagnostics using a TTY terminal as the console display:

- Support for some TTY terminals is optionally installed. If you attach a TTY terminal to a system to run diagnostics, be aware that it may not work properly since AIX support for the terminal may not be installed.
- See the operator manual provided with your terminal to find the key sequences required to respond to the diagnostics.
- When the diagnostics present display information through the S1 port, certain attributes are used. These attributes are set as if the

diagnostics were using an IBM 3161 display terminal. The attributes for the 3151, 3161, and 3164 are listed in the *Diagnostics Information for Multiple Bus Systems*, SA38-0509, Appendix D.

6.2 Standalone Diagnostics

Standalone Diagnostics provide a method of testing the system without running AIX or without having diagnostics installed. This is useful when Online Diagnostics are not available, when devices that are in use during the operation of AIX are to be tested, or if the system is in a state where the Online Diagnostics cannot be run.

The Standalone Diagnostics are packaged on removable media. They must be booted or mounted before they can be run. If booted, you have no access to the AIX error log or to the AIX configuration data, but you can test all the components in your system except for the SCSI adapter used to boot the diagnostics CD-ROM. If the diagnostics are run in concurrent mode using Standalone Diagnostics (the CD-ROM is mounted while AIX is running), then you have access to the AIX error log and the AIX configuration data, but you are limited to the devices you can test.

The following are some considerations that you should make when running Standalone Diagnostics:

- The diagnostics disc must remain in the CD-ROM drive for the entire time that diagnostics are executing. The disc can only be ejected after the system has been turned off and then turned on again (Standalone Mode), or after the diagnostics program has terminated (Online concurrent mode).
- The CD-ROM drive from which diagnostics are loaded cannot be tested.
- The SCSI adapter (or circuitry) controlling the CD-ROM drive from which diagnostic are loaded cannot be tested.

6.2.1 Booting the Standalone Diagnostics CD-ROM

To start the Standalone Diagnostics, perform the following steps:

1. Remove all diskettes and CD-ROMs.
2. Turn the system unit power on.
3. Insert the diagnostics CD-ROM into the CD-ROM drive.
4. When the keyboard POST indicator appears, press **F5** on the directly-attached keyboard, or the number **5** key on the ASCII terminal keyboard, to indicate to the system that the default Boot Device Order List is to be used.

5. Enter any requested password, if set. (Power-On or Privileged-Access).
6. Follow the instructions to select the console.
7. Follow the displayed instructions to test the desired resources.
8. When testing is complete, use the **F3** key to return to the Diagnostic Operation Instructions screen.

6.2.2 Running Standalone Diagnostics in Concurrent Mode

There might be a situation where Online Diagnostics is not available, but you have access to the Standalone Diagnostics CD-ROM. To keep the system available to the users, however, you do not want to reboot and run the diagnostics programs in service mode. In this case, you can mount the CD-ROM and run diagnostics concurrently. For this procedure, perform the following steps:

1. From the command line, type:

```
diag
```

2. If you already have a CD-ROM in the CD-ROM drive, unmount any file systems on the CD-ROM and remove the CD-ROM.
3. You will be asked whether you want to run diagnostics from CD-ROM. Answer *y*.
4. Insert the CD-ROM.
5. Press **Enter** if you want to use the default CD-ROM drive. If you are using a CD-ROM that is not the default, enter its name.

Diagnostics will then mount the CD-ROM, and you can access diagnostics as if they were loaded on your system. When you exit diagnostics from the CD-ROM, the CD-ROM will be unmounted automatically.

6.3 Online Diagnostics

The Online Diagnostics are an AIX Version 4 installable package. They are packaged under `bos.diag` (you will also need to have the `bos.rte.diag` files set installed). There is a separate, installable option for each device or adapter supported by diagnostics. If you install your system from an AIX Version 4 CD-ROM, diagnostics will be automatically installed on your system, including diagnostics support for any PCI devices that are physically installed in your system. In addition, if you install any PCI devices and configure them to the system using `cfmogr -i` (or use **SMIT** -> **Devices** -> **Install/Configure Devices Added After IPL**), specifying an installation device, then diagnostics support for the device (if available) will be automatically installed. Running

Online Diagnostics provides you with access to the AIX error log and the AIX configuration data, but again, the devices you can test are limited.

If the diagnostics package, bos.diag is installed on your system, you can determine which devices in your system have diagnostic support loaded by performing the following steps:

1. Enter `diag`.
2. Press **Enter** on the main screen, and then select **Task Selection**.
3. Select **Display Configuration and Resource List**.

If a device installed in your system is not listed, refer to your installation media to see if the device comes with a diagnostics support fileset.

Online Diagnostics can be run in the following three modes:

- Service mode allows checking of most system resources.
- Concurrent mode allows the normal system functions to continue while selected resources are being checked.
- Single-user mode allows checking of most system resources.

These three modes are discussed in the following sections.

6.3.1 Service Mode

Running diagnostics in service mode provides the most thorough test of the system resources available with Online Diagnostics. As in the concurrent mode, however, there are resources that cannot be tested. These resources are the SCSI adapter and disks used for paging. You can also analyze the error log in this mode by selecting the **Problem Determination** option on the Diagnostic Routine Selection menu.

To access diagnostics from service mode, perform the following steps:

1. Stop all programs including the AIX operating system (`shutdown -F`).
2. Remove all tapes, diskettes, and CD-ROM discs.
3. Turn the system unit off and wait for at least 30 seconds.
4. Turn the system unit on again.
5. After the first POST indicator appears on the system unit console, press **F6** on the directly-attached keyboard, or the number **6** key on the ASCII terminal keyboard, to indicate to the system that the customized Boot Device Order List is to be used. Normally, the first boot device in this list is your boot disk, which should also contain the bos.diag package.

Note

The term POST indicator refers to the icons (graphic display) or device mnemonics (ASCII display) that are displayed while the POST is executing.

6. Enter any requested password, if set (Power-On or Privileged-Access).
7. Select the terminal type (for example, ibm3151 or lft).

Once diagnostics has loaded, you can follow the instructions given on the console to test your system. After you have tested your system, you should exit diagnostics by pressing the **F3** key and follow the instructions displayed to shut down AIX. Failure to shut down AIX could result in damage to data on your disks.

6.3.2 Concurrent Mode

The concurrent mode provides a way to run Online Diagnostics on the system resources while AIX is up and running.

Since the system is running in normal operation, some resources cannot be tested in concurrent mode. The following list shows which resources cannot be tested:

- SCSI adapters used by disks connected to paging devices
- The disk drives used for paging
- Memory
- Processor

There are three levels of testing in concurrent mode:

- The Share-test level tests a resource while it is being shared by programs running in normal operation. This test is mostly limited to normal commands that test for the presence of a device or adapter.
- The Sub-test level tests a portion of a resource while the remaining part of the resource is being used in normal operation. For example, this test could test one port of a multiport device while the other ports are being used in normal operation.
- The Full-test level requires the device not be assigned to, or used by, another operation. This level of test on a disk drive may require the use of the varyoff operation. The diagnostics display menus allow you to vary off the desired resource.

Error Log Analysis can be performed in concurrent mode by selecting the **Problem Determination** option on the Diagnostic Routine Selection menu.

To run Online Diagnostics in concurrent mode, just enter the following command as root:

```
diag
```

After a period of time (up to three minutes), you will be presented with the diagnostics initial screen.

6.3.3 Single-User Mode

Diagnostics can run from single-user mode (maintenance mode – although this is different from booting into maintenance mode from a CD-ROM or a tape). You can shut down your system into a single-user mode that will allow testing of many of your system resources that would normally be in use while AIX is running. All the system resources, except the SCSI adapters and disks used for paging, memory, and processors, can be tested. Error log analysis is available in this mode. Other users will not be allowed to log into the system while the system is in single-user mode.

Logged on as root, you can enter diagnostics in single-user (maintenance) Mode by performing the following:

1. Stop all applications and log all users off the system.
2. Enter `cd /` to change into the root directory.
3. Enter the `shutdown -m` command.
4. When you see a screen similar to this one:

```
Init: Single User Mode
Password:
```

The system prompts you for the root password (if set).

5. Enter the root password.
6. At the command prompt (`#`), enter the `diag` command.
7. When the Diagnostic Operating Instructions area is displayed, follow the instructions to check the desired resources.

After you have completed testing the system, press **F3** to return to the diagnostic operating instructions. Press **F3** again to return to the AIX operating system prompt. Press the **Ctrl-D** key sequence to log off from root.

The following message appears:

```
Do you wish to leave single user mode?  
Enter (Y)es, (N)o or <cr> for the default runlevel:
```

Just press **Enter** to go back to the default runlevel.

6.4 Diagnostic Tasks and Service Aids

The AIX diagnostic package contains programs that are called tasks. This section describes the tasks available in AIX diagnostics Version 4.2 or later (screen output is based on your installed hardware). For more detailed information, refer to *RS/6000 7046 B50 User's Guide*, SA38-0563. The tasks are as follows:

- **Add Resource to Resource List**
Use this task to add resources back to the resource list.
- **Analyze Adapter Internal Log**
The PCI RAID adapter has an internal log that logs information about the adapter and the disk drives attached to the adapter. Whenever data is logged in the internal log, the device driver copies the entries to the AIX system error log and clears the internal log.
- **Backup and Restore Media**
This service aid allows verification of backup media and devices.
- **Certify Media**
This task allows the selection of diskette or hardfiles to be certified.
- **Change Hardware Vital Product Data**
Use this service aid to display the Display/Alter VPD Selection menu. The menu lists all resources installed on the system.
- **Configure Dials and LPFKeys**
This service aid provides a tool for configuring and removing dials/LPFKeys to the asynchronous serial ports.
- **Configure Reboot Policy**
This service aid controls how the system tries to recover from a system crash.

- **Configure Remote Maintenance Policy**
This policy includes modem configurations and phone numbers to use for remote maintenance support.
- **Configure Ring Indicate Power-On Policy**
This service aid allows you to power-on a system by telephone from a remote location (runs only on CHRP/RPA system units).
- **Configure Surveillance Policy**
This task monitors the system for hang conditions, that is, hardware or software failures that cause operating system inactivity.
- **Create Customized Configuration Diskette**
A Standalone Diagnostic Package Configuration Diskette is created in order to change the default refresh rate for a LFT from the console.
- **Delete Resource from Resource List**
Use this task to delete resources from the resource list.
- **Disk Maintenance**
You can make a disk-to-disk copy to another SCSI drive of similar size, as well as display and alter information on disk sectors.
- **Display Configuration and Resource List**
This service aid displays the item header only for all installed resources.
- **Display Firmware Device Node Information**
This task displays the firmware device node information that appears on CHRP/RPA platforms.
- **Display Hardware Error Report**
A tool for viewing the hardware error log is provided.
- **Display Hardware Vital Product Data**
This service aid displays all installed resources, along with any VPD for those resources.
- **Display Microcode Level**
The current level of microcode on an adapter or device can be displayed.
- **Display or Change Bootlist**
This service aid allows the bootlist to be displayed, altered, or erased.

- **Display or Change Diagnostic Run Time Options**
The Display or Change Diagnostic Run Time Options task allows the diagnostic run time options to be set.
- **Display Previous Diagnostic Results**
Each time the diagnostics produce a service request number (SRN) to report a problem, information about that problem is logged. The service representative can look at this log to see which SRNs are recorded.
- **Display Resource Attributes**
This task displays the Customized Device Attributes associated with a selected resource.
- **Display Service Hints**
This service aid reads and displays the information in the CEReadME file from the diagnostics media.
- **Display Software Product Data**
This task invokes SMIT to display information about the installed software.
- **Display System Environmental Sensors**
This service aid displays the environmental sensors implemented on a CHRP/RPA system.
- **Display Test Patterns**
This task provides a means of adjusting system display units by providing displayable test patterns.
- **Download Microcode**
This service aid provides a way to copy microcode to an adapter or device.
- **Flash SK-NET FDDI Firmware**
This task allows the Flash firmware on the SysKonnnect SK-NET FDDI adapter to be updated.
- **Format Media**
The selection of diskettes, hardfiles, or optical media can be formatted.
- **Generic Microcode Download**
This service aid provides a means of executing a generic script from a diskette. The intended purpose for this generic script is to load microcode to a supported resource.

- **Local Area Network Analyzer**
This selection is used to exercise the LAN communications adapters, such as token-ring, Ethernet, and Fiber Distributed Data Interface (FDDI).
- **Periodic Diagnostics**
This selection provides a tool for configuring periodic diagnostics and automatic error log analysis.
- **Run Diagnostics**
The Run Diagnostics task invokes the Resource Selection List menu. When the commit key is pressed, diagnostics are run on all selected resources.
- **Run Error Log Analysis**
The Run Error Log Analysis task also invokes the Resource Selection List menu. When the commit key is pressed, Error Log Analysis is run on all selected resources.
- **Save or Restore Hardware Management Policies**
Use this service aid to save or restore the settings from Ring Indicate Power On Policy, Surveillance Policy, Remote Maintenance Policy, and Reboot Policy.
- **SCSI Bus Analyzer**
This service aid provides a means to diagnose a SCSI bus problem in a free-lance mode.
- **SCSI Device Identification and Removal**
This service aid allows you to choose a SCSI device or location from a menu and to identify a device located in a system unit that uses a SCSI Enclosure Services (SES) backplane.
- **SCSD Tape Drive Service Aid**
This task provides a means to obtain the status or maintenance information from a SCSD tape drive.
- **Spare Sector Availability**
This selection checks the number of spare sectors available on the optical disk.
- **SSA Service Aid**
This service aid provides tools for diagnosing and resolving problems on SSA-attached devices.

- Update Disk-Based Diagnostics
Use this service aid to apply fixes (APARs) on your system.
- Update System or Service Processor Flash
This selection updates the system or service processor flash for CHRP/RPA system units.

6.5 Understanding the Firmware

The firmware is the first code to execute when the machine is powered on. This code resides on the hardware and is loaded at the factory.

The firmware uses the console and the speaker as a means to communicate the results of certain tests. By default, the firmware uses the serial port S1 to display logos, messages, and errors. In case a supported graphics adapter is installed in the machine, the firmware will redirect the output to this adapter instead of the serial port.

6.5.1 General RS/6000 Boot Sequence

RS/6000 CHRP and RPA platforms will display a graphical or text message (hereafter referred to as the boot prompt) after the banner step that identifies the platform and communicates to the user that there are options that may be invoked to alter the boot process.

Once the boot prompt is displayed, the System Management Services (SMS) menu can be invoked (by pressing the **F1** key). SMS provides a user interface for utilities, configuration, and the Multiboot menu for boot/install and the Open Firmware command interpreter.

For boot failures related to firmware trying to access a boot device, it is appropriate for the platform to retry the boot operation, especially in the case of booting from a network device. However, in platforms that have a service processor (for example, Model 260), there are several other types of detected errors for which a reboot retry may be appropriate, for example, check-stops or loss of communication between firmware and the service processor.

The following sections describe the boot phase of operation.

6.5.1.1 Identify and Configure System Components

The firmware must, by various means, become aware of every component in the system and configure or reset those components into a known state. Components include all bridges and device controllers, but may exclude devices that are not involved in the boot process.

Firmware is generally written with a hardware platform in mind so that some components and their configuration data can be hardcoded. Examples of these components are number and type of processors, cache characteristics, and the use of components on the planar. This hardcoding is not a requirement, only a practical approach to a part of this task.

Certain system information must come from walking the I/O buses. This is a technique that will yield identification of I/O device controllers and bridges that reside on modern, well-behaved buses such as PCI.

6.5.1.2 Generate a Device Tree

The firmware will build a device tree. Certain configuration information (configuration variables) may be stored in non-volatile memory.

6.5.1.3 Initialize/Reset System Components

The AIX operating system requires devices to be in a known state at the time control is transferred from the firmware. Firmware may gain control of the hardware in various states depending on what has initiated the boot process.

6.5.1.4 Locate an OS Boot Image

A device and filename can be specified directly from the AIX boot command. Once the boot image is located, the device path is set in the device tree as the bootpath property of the chosen node.

If multi-boot (multiple bootable operating systems residing on the same hardware platform) is supported, a configuration variable instructs the firmware to display a multi-boot menu from which the OS and bootpath are selected.

6.5.1.5 Load the Boot Image into Memory

After locating the image, it is loaded into memory at the location given by a configuration variable or as specified by the OS load image format.

6.5.2 Typical Boot Sequence for RS/6000 7046 Model B50

After the system unit power is turned on, the system firmware tests the system memory and identifies failing memory cards and memory module locations. System firmware checks to see what devices are in the system and performs a simple test on them. The system firmware displays the device name or device icon being tested. After the keyboard name or icon appears, you can enter the Systems Management Services menu by pressing the **1** key (on a ASCII terminal) or the **F1** key (on a graphics terminal).

You can also use one of the following:

- **5** (on a ASCII terminal) or **F5** (on a graphics terminal) to start the Standalone Diagnostics (CD).
- **6** (on a ASCII terminal) or **F6** to start the Online Diagnostics (hard disk)

Depending on the boot device, a checkpoint may be displayed on the operator panel for an extended period of time while the boot image is retrieved from the device. If a bootable image is found, the system unit loads and starts the operating system. The AIX boot code indicates LED progress codes in the range from 05XX to 09XX. After the AIX Boot is completed, the AIX login prompt appears on the main console.

Note

For more information about the progress codes, refer to *Diagnostics Information for Multiple Bus Systems*, SA38-0509. Be sure to read the chapter on Diagnostics Numbers and Codes.

If no bootable operating system image is found on any of the storage devices listed in the system firmware boot list, then the system unit attempts to boot again. With each attempt, you have the opportunity to start the System Management Services (SMS) as described in 6.7.1, “SMS Startup” on page 117.

6.6 Power-On Self-Test

Power-On Self-Test (POST) is the process by which the firmware tests those areas of the hardware that are critical to its ability to carry out the boot process. It is not intended to be all-inclusive or sophisticated as to how it communicates with the user. Diagnostics with these characteristics will generally be provided as a service aid. The following sections describe power-on self-tests that are performed by the firmware on PCI-based RS/6000 systems.

6.6.1 Processor POST

The firmware tests the processor very briefly. If there is a failure, the system will stop, and the initial beep of the speaker will not be heard.

6.6.2 Memory DIMM and L2 Cache POST

The firmware RAM POST tests all the memory. This is done in segments. After the initial short beep, if the system stops with the hard file indicator on

solid, this is the indication that no memory was found. If some memory is found, then a working 2 KB continuous segment is sought. If the 2 KB segment cannot be found, the system stops with the hard file indicator on solid.

The L2 cache is tested. If it fails, there will be a momentary indication when the display becomes available. The error is also written to the error log. During the system initialization, when icons are being displayed for the individual subsystems, one of the icons will be the memory icon. During the display of this icon, an 8-digit error code will be displayed if the L2 cache is not fully functional.

The remainder of memory is also tested, and if a failure is detected, an 8-digit error code is posted. This error indicates which Dual Inline Memory Module (DIMM) slot is at fault. The basic algorithm for the system memory configuration is as follows:

1. Query each memory slot to see if there is a memory DIMM installed.
 - If yes, determine its size. Identify each DIMM as IBM or non-IBM, ECC, or parity. This is done by querying special identification registers in the DIMMs themselves.
 - If the DIMM could not be identified, display the memory slot number, as well as a warning saying that a DIMM could not be identified. Proceed with the next DIMM.
2. Test the DIMM.
 - If this is an IBM DIMM, and it passes the test, make it available and add its size to the available memory quantity. If it does not pass the test, keep the DIMM out of the system configuration and display a warning on the screen reporting the faulty DIMM slot number.
 - If this is a non-IBM DIMM and passes the test, make it available and add its size to the available memory quantity. In addition, display a message on the screen reporting that a non-supported, non-IBM memory DIMM has been found, tested, and correctly configured. If the DIMM does not pass the test, display a warning on the screen reporting that a non-IBM memory module was found and failed to pass the test.
3. Proceed with the next DIMM.

At the end of this process, the system will have available all the memory it could configure (IBM branded or not).

6.6.3 Graphics Adapter POST

The firmware tests the graphics adapter but not the display. The system will give a speaker beep sequence when the graphics adapter is found to be non-functional. The beep sequence is one long beep, followed by one short beep. The error is logged, and the system boot continues.

6.6.4 Keyboard and Mouse Controller POST

The keyboard/mouse controller has a built-in Power-On Self-Test, and the firmware checks the result. If an error is detected, an 8-digit error code is displayed. The error is written to the error log, and the system continues.

6.6.5 SCSI Subsystem POST

The SCSI POST opens the adapter and runs several test scripts that check and exercise the various functions related to the chips on the integrated SCSI adapter, such as DMA, simulated mailbox commands, and so on. If an error is detected, an 8-digit error code is displayed, the error is written to the error log, and the system continues.

Note

The 8-digit POST error codes are listed and described in the Service Guide for the corresponding PCI-based RS/6000 system:

RS/6000 7046 B50 Service Guide, SA38-0564

6.7 System Management Services (SMS)

The SMS programs make it possible for you to view information about your computer, run memory tests, and set the device from which the server will boot.

By default, the firmware uses the serial port S1 to display logos, messages, errors, and also as the interface to the SMS programs. If a supported graphics adapter is installed in the machine, the firmware will redirect the output to this adapter instead of the serial port. This means, if you have installed a graphics adapter, but you want to use an ASCII terminal attached to serial port S1 as the console for the SMS programs, you will have to remove the graphics adapter first.

The SMS programs provide two types of interfaces: a Graphical User Interface (GUI) that can be used on a graphical display and an ASCII

interface that can be used either on an ASCII terminal or on a graphical display.

Note

The ASCII interface provides more functions than the graphical interface.

To start the programs:

1. Turn on or shut down and restart your computer.
2. When the first screen appears, press the **F1** key to use the GUI. On an ASCII terminal, press the number **1** key.

Note

You must press keys listed above before the last indicator appears.

6.7.1 SMS Startup

During the boot process, function keys can be pressed to alter the electronic key-switch and determine which boot list to use during the boot procedure. There are four function keys that determine the type of boot that will occur. These are described as follows:

F1 key (1 with an ASCII terminal)

The **F1** key pressed on a directly-attached keyboard will invoke the SMS GUI interface. Press the **1** key at an ASCII terminal and it will invoke the SMS ASCII interface. The electronic key-switch will be set to Normal Mode.

Mksysb Booting

Booting from a tape (to restore a mksysb tape or boot into maintenance mode) requires SMS. From SMS, you can select an alternative boot device that includes a tape drive if one is installed.

F5 key (5 with an ASCII terminal)

The **F5** key will cause the default firmware boot list to be used. The electronic key-switch is set to service mode. This mode attempts to boot from the first device of each found in the list. It does not search for other bootable devices of that type if the first device is not bootable. Instead, it continues to the next device type in the list. If, after one pass through the boot list, no boot record is found, the firmware will invoke the SMS program.

F6 key (6 with an ASCII terminal)

When **F6** is selected, the firmware looks for a boot record according to the custom boot list that was set up by SMS or by using the `bootlist` command. If the boot list was not customized, the default boot list is used. The electronic key-switch is set to service mode. Again, if, after a pass through the custom boot list, no boot record is found, the firmware will invoke the SMS program.

F8 key (8 with an ASCII terminal)

Press **F8** to access the Open Firmware `OK>` prompt. This has no influence on the selection of the boot list or key-switch.

Note

When AIX is installed on a machine, it automatically sets the customized boot list to the installed hard disk. The customized boot list is thereafter used by the system when booting in normal mode.

6.7.2 SMS Graphical Main Menu

The graphical System Management Services main menu offers the following tasks, as shown in Figure 20 on page 119:

- | | |
|------------------|---|
| Config | Enables you to view your system setup. |
| Multiboot | Enables you to set and view the default operating system, modify the boot sequence, access the Open Firmware command prompt, and other options. |
| Utilities | Provides several tools to set and remove passwords, to set the unattended start mode, to update the firmware, to set the IP addresses for network booting, and other utilities. |



Figure 20. System Management Services - Main Menu

6.7.2.1 SMS Submenu Config

Selecting this icon makes it possible for you to view information about the setup of your system unit. A list similar to the following appears when you select the Config icon. If more than one screen of information is available, a blue arrow in the top right corner of the screen appears. Use the **Page Up** and **Page Down** keys to scroll through the pages.

6.7.2.2 SMS Submenu Multiboot

The options, as shown in Figure 21 on page 120, available from this screen allow you to view and set various options regarding the operating system and boot sequence:

- Select Software** The Select Software option allows you to choose which operating system to use.
- Software Default** This option enables you to select an operating system as the default operating system for your system unit.
- Install From** Enables you to select a media drive from which to install an operating system.

- Boot Sequence** Enables you to view and change the custom boot list (the sequence in which devices are searched for operating system code).
- OK Prompt** Enables you to go directly to the Open Firmware command prompt. The Open Firmware command prompt is used for debugging and device driver development.
- MultiBoot Startup** Clicking this button toggles whether the Multiboot menu will appear automatically at startup.

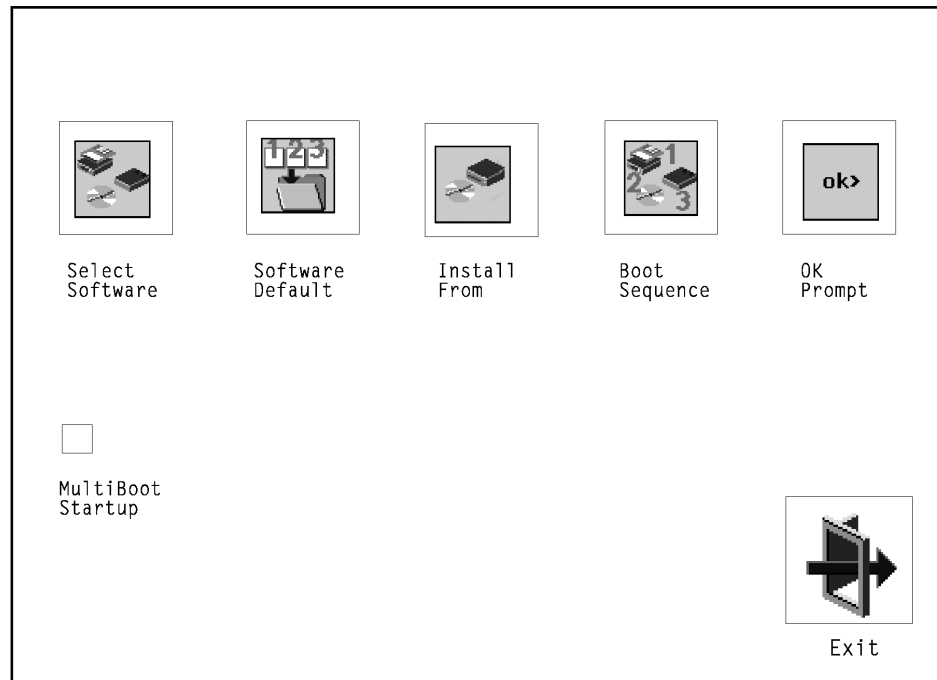


Figure 21. SMS Submenu Multiboot

6.7.2.3 SMS Submenu Utilities

Selecting this icon enables you to perform various tasks and view additional information about your system unit, as shown in Figure 22 on page 121.

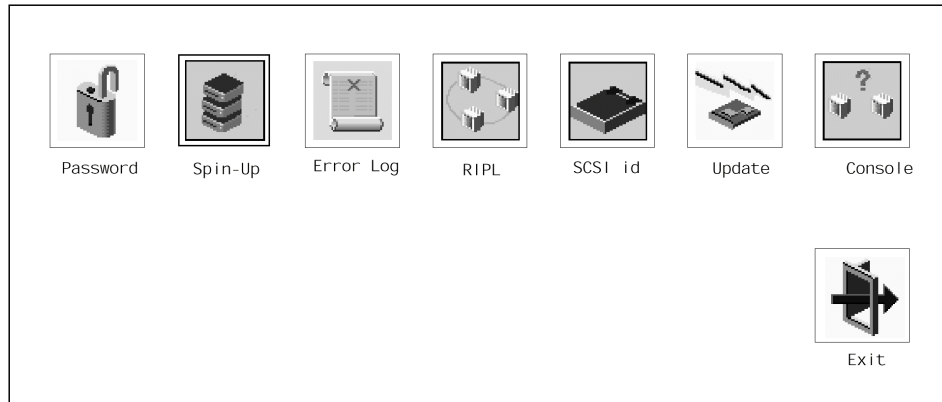


Figure 22. SMS Submenu Utilities

The following icons are described in the order that they appear:

- Firmware Passwords

You can limit access to the system by using passwords. The passwords can be up to eight characters long and are stored in NVRAM. There are three categories of restricted access:

- Power-On Password

When the system is powered on, it checks whether a Power-On Password (POP) is present. If there is one present, and the unattended start mode is not set, it means the machine's owner does not want the system to be used unless the POP password is supplied. In this case, the system will prompt for the POP password. You are given three attempts to enter the correct password. If you fail to supply the correct password, the system will go into a hung state and must be powered-off before continuing.

- Privileged Access Password

The Privileged Access Password (PAP) is used to limit access to the System Management Services (SMS) programs (also called the Supervisory Password). When you press one of the keys to access SMS, the system will check to see if a PAP exists, and if it does, you are prompted to enter the PAP. You are given three attempts to supply the correct password. If you fail to do so, the system will go into a hung state and must be powered-off before continuing.

- Unattended Start Mode

To use this mode, a POP must be previously specified. If unattended start mode is enabled, the system will boot from the defined boot device without requiring the you to enter the power-on password. While the system can be booted without entering the POP password, the keyboard controller will be programmed to lock up until the POP password is provided. This mode is ideal for servers that run unattended. After an electrical power failure, for example, the operating system will be rebooted without waiting for you to enter the power-on password.

Figure 23 shows the SMS submenu where you can set the password.

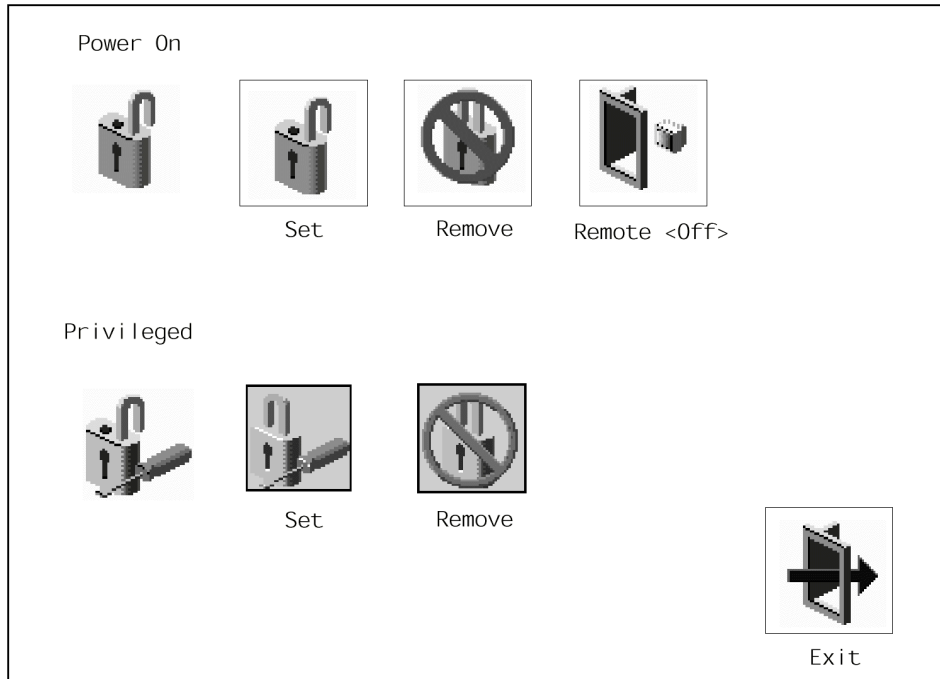


Figure 23. SMS Submenu - Firmware Passwords

Note

For details about enabling and disabling the power-on password and privileged access password, refer to the relevant chapter in the appropriate service guide:

RS/6000 7046 B50 Service Guide, SA38-0564

- **Hard Disk Spin-Up Delay**
Enables you to change the spin-up delay for SCSI hard disk drives attached to your system. Spin-up delay values can be entered manually, or else a default setting can be used. All values are measured in seconds. The default is two seconds.
- **Error Log**
The error log enables you to view and clear the firmware error log information for your system unit.
- **RIPL (Remote Initial Program Load)**
Enables you to select a remote system from which to load programs via a network adapter when your system unit is first turned on. This option also allows you to configure network adapters that require setup. You can ping a remote adapter to check if the network hardware is working correctly.
- **SCSI ID**
Allows you to view and change the addresses (IDs) of the SCSI controllers attached to your system unit.
- **Update**
Allows you to update the system firmware program. Refer to 6.7.4, “SMS Firmware Update” on page 124.
- **Console**
If you have more than one keyboard and display attached to your system unit, or if you have an ASCII terminal attached to your system unit in addition to a keyboard and display, this tool allows you to define which one is active.

6.7.3 SMS ASCII Main Menu

The ASCII version of the System Management Services includes an additional function in the main menu not found in the graphical version. The following languages are supported:

- English

- French
- German
- Italian
- Spanish
- Swedish

All menus and options, as well as the button titles, have been translated. Your ACSII terminal must support the ISO8859-1 character set in order to properly display languages other than English.

6.7.3.1 ACSII Terminal Settings

If you have a graphics adapter in your system, then System Management Services will only display information using that adapter. If you have an ASCII terminal attached and have a graphics adapter installed, nothing will be displayed on the ASCII terminal during the firmware boot process. To be able to use the ASCII terminal to access System Management Services, you will have to physically remove the graphics adapter from the system. If there are no terminals or graphics adapters attached, the system will still be able to boot up.

If you are using an ASCII terminal, the following parameters should be used:

- 9600 baud
- 8-bit
- No parity
- 1 stop bit
- 24 lines by 80 columns
- Echo mode on

6.7.4 SMS Firmware Update

There are several situations where a firmware update may be required. For example, an update may be required after a field Engineering Change (EC) to provide firmware support for new adapters and peripherals or to enhance functions for existing peripherals.

6.7.4.1 How to Determine the Current SMS Firmware Version

The firmware version is displayed on the initial startup screen on a graphical display and at the top of each System Management Services (SMS) menu on an ASCII Terminal display. It can also be found on the configuration display under the SMS menus.

To access the SMS menus, power-on the system if it is off, or shut down and restart the system if it is on. On a graphical display, the firmware version is displayed in the lower right corner of the initial display. Press the **F1** key while the startup icons are displayed to stop at the SMS menus.

On an ASCII Terminal display, press the **1** key at the initial display when the ASCII icons are displayed at the bottom of the screen. Note the firmware version at the top of the menus.

Select **Configuration** from the initial SMS menu to see the system's configuration. The firmware version is near the bottom of this list.

6.7.4.2 How to Obtain the SMS Firmware Update

In order to create a firmware diskette with the latest level of SMS firmware available, download the newest firmware for your machine from:

<http://www.rs6000.ibm.com/support/micro>

There are two types of archives currently being distributed: one for AIX and one for DOS-compatible operating systems, such as DOS, OS/2, and Windows. Both archives are self-extracting executables that will unpack themselves when run.

To unpack the AIX archive, filename.bin, execute the following commands:

```
chmod +x filename.bin
./filename.bin
```

To unpack the DOS-compatible archive, filename.exe, execute the archive:

```
filename1
```

(¹This depends on the model for which you want to download the firmware.)

6.7.4.3 How to Create the SMS Firmware Diskette

The method of creating the firmware update diskette will vary depending on the operating system being used and the distribution archive. In all cases, 1.44 MB DOS-formatted diskettes should be used. For AIX users, DD images should be provided. To create a firmware update diskette, insert a blank diskette into the drive of the system being used. If a makefw script is provided, execute it. Otherwise, type the following command:

```
dd if=filename.dd of=/dev/rfd0
```

For users of DOS-compatible operating systems, discrete files will be provided with a batch file named makefw to copy the appropriate files to diskette. Run the batch file to create the diskette:

6.7.4.4 How to Perform the SMS Firmware Update

Once you have a firmware update diskette available, you are ready to perform the update. Follow the steps below carefully.

The procedure below will guide you through *soft-loading* the new firmware which loads it into memory and runs it. This is a very important step that will ensure that your new firmware is not corrupted and will run on your system as it is configured *before* the new firmware is permanently flashed into the system. If the soft-load fails, *do not* flash the new firmware.

1. Power-on the system, if off, or shut down and restart.
2. Go to the open firmware ok prompt:
 - From a graphics display, press **F8** when the icons appear at startup.
 - From an ASCII terminal, press **8** when the ASCII icons appear during startup.
3. Insert the diskette with the firmware image and soft-load it by typing the following commands at the ok prompt:
 - `> setenv load-base 600000`
 - `> boot floppy:1,\name_of_firmware_level.img`

Note

`name_of_firmware_level` depends on the system for which you want to update the firmware (for example, `tig98153`).

A console selection may be run at this time to allow you to choose a display. A number will appear on each display that can be used for the firmware icons and SMS menus. Press the number of the display you wish to use. A graphics display is preferable.

4. Go to the SMS menu:
 - From a graphics display, press **F1** when the icons appear at startup.
 - From an ASCII terminal, press **1** when the ASCII icons appear during startup.
5. Verify that the new firmware is currently running. See 6.7.4.1, “How to Determine the Current SMS Firmware Version” on page 124. If the new firmware is not running, something went wrong. Go back to step one and try again. If the new firmware *is* running correctly, continue to the next steps to permanently flash the new firmware image.

Note

If power is lost during the firmware update procedure, the system could be rendered inoperable. To avoid this, *do not power off* the system during the firmware update and take any steps to ensure that power is not interrupted during the update procedure.

6. From a graphics display, use the firmware update utility. Select **Utilities** from the menu, then **Update** and follow the on-screen instructions to complete the update procedure.
7. From an ASCII terminal, use the firmware update utility. Select **Utilities** from the main menu, then **Update system firmware** and follow the on-screen instructions to complete the procedure.

6.7.4.5 SMS Firmware Recovery

If the system will not boot up fully, and the checkpoint display reads FEB, the firmware is partially corrupted but should be recoverable. Make a copy of the firmware update diskette, rename the .img file on this diskette to PRECOVER.IMG, and insert this recovery diskette in the drive when FEB appears on the checkpoint display.

If the system drops to the ok prompt at checkpoint F4D, this may also be a partial corruption of the firmware. Use the ok prompt to attempt the soft-load procedure and re-flash the new firmware. Start at step three under 6.7.4.4, “How to Perform the SMS Firmware Update” on page 126.

If the checkpoint display reads 888 when the power button is pushed, the system firmware is completely corrupted. Report the problem to your service organization.

Note

The diskette containing the new firmware update will be your firmware recovery diskette.

6.7.5 SMS Firmware Checkpoints

Firmware Checkpoints are intended to let users and service personal know what the system is doing, in some detail, as it initializes. These checkpoints occur during system startup and are not intended to be error indicators, but in some cases, a system could hang at one of the checkpoints. The most appropriate action is included with each checkpoint.

6.7.5.1 RS/6000 7046 Model B50 Firmware Checkpoints

The firmware for the Model B50 uses progress codes (checkpoints) in the range of E1xx to EFFF. The Exxx-Codes are documented in the *RS/6000 7046 B50 Service Guide*, SA38-0564.

Chapter 7. Application Software Solutions for the Model B50

In this chapter, the basic software solutions applicable for the Model B50 are discussed.

An Internet Service Provider (ISP) is a business that has access to the Internet and sells this ability to connect to the Internet to members of the general public. An Internet Service Provider is not the same as an Information Service or contents provider. Recently it was not easy to distinguish between an Internet Service Provider and a contents provider and any other network-related service provider. These providers offered not only access to Internet but also offered additional content and specialized services for their customers. Online periodicals, special chat rooms, stock quotes, and educational material are some examples of additional content. However, these kinds of information services are becoming more and more entwined with the Internet and now almost all provide the ability to directly access the Internet.

An ISP provides various basic services for customers in addition to connectivity to the Internet, such as mail transfer, name service, Web hosting, and have a firewall to keep their services safe from the rest of the Internet.

For Service Providers, the evolution from business to e-business provides significant new opportunities to grow revenue. As businesses move to the Web, they're looking for value-added services that make them more successful. Many naturally turn to Service Providers, for answers.

IBM has created a series of bundled hardware and software packages for essential service areas: caching, security, systems management, Customer Relationship Management (CRM), commerce, messaging, and hosting. These packages deliver assured, scalable solutions that enable you to help your customers implement e-business.

Moreover, packages can be mixed and matched. You can select only the solution and performance level you need, from low-cost entry systems to high-value, high-performance premium configurations.

Table 28 on page 130 lists sample application packages now available in the key areas of messaging, hosting, commerce, caching, security, network, and

subscriber management. Many additional solutions from independent software vendors run on reliable, scalable, RS/6000 servers.

Table 28. Selected Service Provider Offerings

	Entry	Intermediate	Advanced
Messaging	InterMail Post.Office	InterMail KX	Intermail MX Integrated Messaging
Hosting	WebSphere Application Server	WebSphere Application Server—Advanced	WebSphere Application Server—Enterprise SP WebSphere Performance Pack
Commerce	Net.Commerce Start	Net.Commerce CHS	Net.Commerce CHS IBM Payment Server
Caching	WebSphere Performance Pack Cache Manager	WebSphere Performance Pack	WebSphere Performance Pack Distributed Caching
Security	SecureWay Firewall	SecureWay Software for VPNs	First Secure Managed Security Services
Management	Tivoli Director	Tivoli TME	Tivoli TME Tivoli Cross Site
Subscriber Management		Intelligent Subscriber Management System (ISMS)	ISMS

The following sections provide an overview of the basic services ISPs use together with a small sample of the LPPs that IBM provides for this industry are discussed.

7.1 Internet Services for ISP

There are several services you should consider supporting for your user base. This section outlines several of the key services commonly supported by ISPs. It is important to note that you are not expected to run a server for every single service discussed here. You should treat this list as food for thought. You may also find that some, or all, of these services may be

provided either free (included in the cost of your link), or at an additional cost from your upstream provider. Throughout this section, server refers to the program running on one of your machines providing the service being discussed. In most cases, you will be able to run more than one server on each machine.

7.1.1 Domain Name Service

The Domain Name Service (DNS) has become the glue that binds the Internet together. It provides a mechanism for converting easy-to-remember names, such as `www.ibm.com`, into the less easy to remember IP addresses that are used in the underlying protocols. It is also used for other services, for example, using a special record in the DNS. You can make use of your upstream provider's mail backup servers (if they provide that service).

7.1.1.1 Berkeley Internet Name Domain

Before you can register any domains, you need to have the domains configured on a name server. If you choose to run your own name server, the most commonly used server is Berkeley Internet Name Domain (BIND, which is now maintained by the Internet Software Consortium (ISC). Other DNS implementations have been made available, but the majority of name servers in the field are either running BIND, or a product that is based on BIND. BIND is released in source code format for free by the ISC, and a lot of effort has been made to support as many operating systems as possible. If you are running UNIX as your server platform, the chances are that the provided DNS daemon is an implementation of BIND. The support Web page for BIND can be found at <http://www.isc.org/bind.html> and it includes many links to other DNS-related sites. BIND has its own support newsgroup at `comp.protocols.dns.bind`. The latest version (as of August 1999) of BIND is 8.2.1.

7.1.2 Mail Service

It used to be the case that if you provided an e-mail address for your users, you were classed as an ISP. Although this perception has changed, e-mail is still a critical service to provide. Your users will expect at least one e-mail address from you, and most ISPs now provide around three e-mail addresses per account. You will need two mail servers: one for your users to collect their own mail (POP server) and one to receive the incoming mail and place it on the POP server and allow your users to send mail (SMTP server or relay).

7.1.2.1 POP/IMAP4 Server

Because your dial up users are not typically connected to the Internet 24 hours a day, they are not usually connected when somebody sends them mail. Therefore, you will have to hold their mail for them until they pick it up.

The most common method of mail retrieval by clients is via the Post Office Protocol Version 3 (POP3). The user's e-mail software connects to the POP server, logs on with a user ID and password, downloads any waiting mail, deletes the mail from the server, and disconnects. Most UNIX operating systems come with a POP server supplied, but there are several alternatives available on the Internet.

Internet Mail Application Protocol, currently at Version 4 (IMAP4), is less common than POP3 but is gaining popularity all the time. The most significant difference between POP and IMAP is that IMAP clients leave the mail on the server, rather than downloading the messages and removing them from the server as POP clients do. IMAP provides folders on the server to provide a remote mailbox which can be manipulated in the same way as local mailboxes.

7.1.2.2 SMTP Server

The way that e-mail is sent from source to destination has changed very little since it was first used. It used to be the case that the source machine connected directly to the target machine, transferred the notes, and disconnected. If the target machine was down, the source machine would try again later, and keep trying until either the mail was delivered, or some time-out limit was reached. However, some machines needed to receive e-mail but were not directly connected to the Internet. This was accomplished by placing mail relays on the Internet that knew how to contact these non-Internet connected machines.

These principles still hold, but the mail relays now have an extra role to perform since some, or all of your customers, are not connected to the Internet 24 hours a day. So if the destination is down, their machines may not be able to retry. The solution is to use a mail relay. In this case, the user's e-mail software sends the mail to your mail relay, which then attempts to send it on to the destination on behalf of the user.

Every UNIX implementation comes with a mail server. The most popular one is Sendmail (<http://www.sendmail.org/>), which is supported by its author, Eric Allman. Sendmail has now been implemented as a commercial product although open source sendmail is also available.

7.1.3 Web Service

The World Wide Web (WWW) is the most popular service on the Internet and one of common applications in ISP services. In today's Internet, you are nobody without a Web site. Your users will expect some space on your Web server to put up some pages of their own. This can be accomplished by either asking your users to e-mail you their Web pages and graphics for you to upload onto the Web server, or by giving each user upload access to their own area on the Web server. To host a Web server, you should use an HTTP server, such as the IBM HTTP Server, and manage the contents and security.

7.1.4 FTP Service

File Transfer Protocol (FTP) is a simple protocol that is supported by all Internet server and client platforms. An FTP server can be used to distribute updates to client programs to your users, and your users may want to share data with other people by using FTP.

7.1.5 Directory Service

Directory services are becoming more important because the Internet name space is becoming difficult to manage. LDAP services are a growing part of ISP services.

7.1.6 Security and System Management

The widespread Internet connectivity and penetration to common company activities requires more security and system management. Creating and maintaining firewalls and keeping the system free of computer viruses are important issues for ISPs.

7.2 Applications for Businesses

The services on the Internet are commonly duplicated inside companies for business intranet activities. In such environments, some services are required in addition to the basic services. Business users need file sharing, printer sharing, and advanced functional messaging systems, such as Lotus Notes. These services are only introduced here since a detailed discussion is beyond the scope of this publication.

7.2.1 File Sharing Service

Interoperability and transparent access to files either on UNIX or on Windows is a requirement today. To access files, each platform must be able to

understand the same protocol and services, such as TCP/IP, NetBIOS and CIFS, SMB, and NFS to name a few.

7.2.2 Printer Sharing Service

Printer sharing allows the optimal placement of printers and the distribution of print services with respect to users, independent of the client operating system.

7.3 Applications for Internet Service

IBM offers several software solutions for ISPs or small business users. In this section, we introduce AIX functions and additional LPPs designed for Internetworking.

7.3.1 Basic Internet Features on AIX Version 4.3.3

The AIX Version 4.3 operating system contains basic Internet services as follows. Some new features are introduced with AIX Version 4.3.3.

- POP3 / IMAP Server

The AIX mail server includes Post Office Protocol Version 3 (POP3) and Internet Message Access Protocol Version 4 (IMAP4). POP3 conforms to RFC 1725, and IMAP4 conforms to RFC 1730. The POP server performs Message Transfer Agent (MTA) services for a disconnected enduser. The POP server holds messages, and when the POP client (mail user agent) connects and requests messages, the POP server downloads all of the pending messages to the client machine. Most of the mail clients today operate using the POP protocol. An IMAP server performs online and disconnected access. Additional functions include:

- The ability to append messages to a remote folder
 - The ability to set standard and user-defined message status flags
 - Support for simultaneous update and update discovery in shared folders
 - New mail notification
- Sendmail Version 8.9

Sendmail is upgraded to Version 8.9.3 on AIX Version 4.3.3 providing an anti-spamming feature as well as other 8.9.3 enhancements.

- FTP Server

Standard ftpd is included.

- NIS+

AIX Version 4.3.3 includes a port of the Sun Solaris Version 2.5 NIS+ implementation. This function is provided in addition to the NIS support which is still available. This new naming service provides enhanced capabilities for security management in a distributed system environment, as well as supporting the largest enterprises.

- IPv4 and IPv6

AIX Version 4.3.3 supports IPv4 and IPv6 gateway capabilities in addition to the IPv6 host capabilities supported in AIX Version 4.3. This includes Merit GateD Version 6.0 (named gated) that provides gateway routing functions for the RIP, EGP, BGP, HELLO, OSPF, and SNMP protocols and is limited to unicast routing. Specifically, the AIX Version 4.3.3 version of gated supports intradomain routing using RIP Versions 1 and 2, OSPF Version 2, ISIS, HELLO, and SLSP. GateD Version 6.0 provides interdomain routing using BGP Versions 2, 3, and 4 and EGP. SNMP support, provided through the SMUX protocol, is based on ISODE Version 8. This version of gated has been modified to support IPv6.

- Java Support

The IBM AIX Developer Kit, Java Technology Edition, Version 1.1.8 delivered on AIX Version 4.3.3 provides the following features:

- Java Remote Method Invocation
- Java Database Connectivities (Jdbc)
- Swing 1.1.0
- Big Decimal

- IPSEC/VPN

The IBM SecureWay Communications Server for AIX connects customers, partners, and employees nearby or around the world to resources across multiple systems with secure communications. Both data and transactions are protected from unauthorized access or intrusion through the firewall and use of Virtual Private Network (VPN) technology.

- LDAP Version 2, Version 3

The IBM SecureWay Directory is an open, cross platform server optimized to support Lightweight Directory Access Protocol-enabled (LDAP) applications that integrate enterprise systems. Providing a unified architecture that allows users to share data with people, applications, and network resources, SecureWay Directory helps improve communication, speed development and deployment of Web applications, and increase the security of the network. Utilizing the power of IBM DB2 Universal

Database and its transactional data store, the directory extends the performance and availability of DB2 to an enterprise directory service.

- DNS

BIND 8.1.2 is included in AIX Version 4.3.3.

- RSVP

Advanced Internet services are based on Quality of Service (QoS) and policy-based networking. QoS embodies the idea of preferential treatment for certain traffic flows and enables service providers to implement and effectively manage Service Level Agreements (SLA). The AIX QoS implementation is based on the Internet Engineering Task Force (IETF) standards, Integrated Services (IntServ), and Differentiated Services (DiffServ). IntServ utilizes the Resource ReSerVation Protocol (RSVP) available to applications via the RSVP API (RAPI). DiffServ support includes IP packet marking for IP packets selected via filtering. The AIX QoS also offers bandwidth management functions, such as Traffic Shaping and Policing. Policy-based networking is an industry effort to define and standardize various network configuration parameters needed for effective network management. The AIX QoS scope covers both QoS and policy-based networking. This enhancement to AIX provides system administrators with the benefits of both QoS support and policy-based networking in meeting the challenges of QoS offerings across complex networks.

- SOCKS5 Library

The AIX SOCKS API allows generic TCP/IP applications to connect to hosts through a generic TCP/IP proxy using SOCKS protocol Version 5.

7.3.2 Service Providers Utility Pack

The AIX Service Provider Utility Pack offering that, when combined with AIX Version 4.3 and applications on the Bonus Pack, helps reduce total RS/6000 system set-up time and operational costs for Service Providers. This offering combines the Internet News Server, an Internet FTP Server, Internet Administration, Internet Services Monitor, and Security Hardening.

Base AIX Version 4.3 includes robust scalability and leadership Lightweight Directory Access Protocol (LDAP) support, while the accompanying AIX Version 4.3 contributes the high-performance IBM HTTP Server and WebSphere Application Server at no additional charge.

Find out more about the Service Providers Utility Pack at:

<http://www.alphaworks.ibm.com/tech/spup>

Table 29 lists the files available for download that comprise this offering.

Table 29. Service Providers Utility Pack Components

File	Size	Comments
libsrcsec.so.zip	1 KB	Updated libsrcsec.so to enable encrypted communication between servers
spup.base	5540 KB	Service Provider Utility Pack installable software image for AIX Version 4.3.3
spup.html.en_US	928 KB	Service Provider Utility Pack installable documentation for AIX Version 4.3.3
requisites.tar	1780 KB	This file contains the following requisites that are to be installed on AIX 4.3.3 before this product is installed. bos.rte.SRC.4.3.3.1.bff bos.sysmgt.nim.client.4.3.3.1.bff bos.sysmgt.nim.master.4.3.3.1.bff efix.tar

The AIX Service Provider Utility Pack was built to meet the needs of Internet Service Providers (ISP) and Application Service Providers (ASP) by providing additional key function in a number of areas.

7.3.2.1 Additional Internet Services

Offer additional functionality to the standard FTP and Inews server capability in the AIX operating system.

This NEWS server is derived from InterNetNews (INN). The configuration and administration interfaces are preserved from INN. The news reading can be restricted by user and host and is handled by a separate daemon. The NEWS server can simultaneously accept data from multiple sources and feed multiple peers. The newsgroups accepted can be restricted by peer. Administrators have the capability to write PERL filters to reject articles based on content. In addition, the administrators will have the capability to setup schedules to expire old articles and feed new articles to other sites.

This server can be configured, started and stopped through the Web based Internet Services Management capability.

This FTP server is derived from Washington University-FTPD (WU-FTPD) which preserves the configuration and administration interfaces. Unlike the FTP Server that is part of AIX, the WU-FTPD does not require the inetd daemon running and can be executed as an independent program. In addition to supporting the FTP protocol defined in RFC 959, the WU-FTP has many advanced features and functions which help Service Providers to

control resource utilization, manage user access, and track data transfer and command execution.

This FTP server can be configured and administrated through the Web-based Internet Services Management capability.

7.3.2.2 Internet Services Management and Administration

This capability provides the ability to manage the key internet services (News, FTP and HTTP Server) through a centralized, Web-based interface. It allows users to securely administer these services from any location. In addition, administrators may use the administration console to launch additional, user-defined applications. The distributed management architecture allows geographic and topological separations of administration clients, administration servers, and managed servers, giving users the ability to manage multiple service installations from a single console remotely. Furthermore, these services utilize a consistent look and feel through the Graphical User Interface management, simplifying the process of managing multiple, disparate services.

Primary authentication information is established at the centralized administration server and forwarded to the appropriate service running on managed servers, freeing the administrators from repeated authentication processes. This single sign-on (using centralized administration server) is provided for key internet services as well as user-defined applications (managed through this framework). Alternatively user-defined applications may require additional authentication information. Industry-standard SSL is used for secure connections both from the client browser to the administration server, and from the administration server to the managed servers.

Each service (on each managed server) has distinct authentication controls as well, which gives an administrator the ability to delegate management duties for particular services to alternate users without compromising the security of other services. All administrative actions are logged through AIX syslog, allowing an administrator to monitor the activity of alternate administrative users.

After receiving and authenticating each administration request, the service agent on the server to be managed performs the requested action through modification of the configuration source (files, ODM databases, and so on) that is native to the managed component or application. This allows users who prefer direct editing of component configuration (or third-party configuration tools) to act in concert with the centralized administration interface, rather than having to relearn new file or database formats.

7.3.2.3 Internet Services Monitor

Internet Service Monitor provides the capability to monitor the performance of Internet Services such as HTTP, Mail, News, and FTP services, allowing administrators to more effectively manage load distribution and proactively perform problem determination. Each service is monitored through the use of actual service requests, providing a view of service performance that most accurately simulates actual customer experience. In addition to monitoring the customer-view performance of primary services (HTTP, FTP, NNTP, and POP3), the services monitor also provides a view of server reachability, and can be extended to include monitoring user-defined data.

Users can also request notifications if the performance of a particular service drops below a pre-configured threshold.

7.3.2.4 Security Hardening

Security Hardening capability takes a number of actions to toughen the security of the AIX operating system in the Internet Services environment. The major actions are as follows.

- It configures auditing to monitor administrative actions. Auditing also allows monitoring the access of administrative accounts from non-privileged users. Potential *system crackers* are denied valuable information by turning off non-essential facilities in the `inetd.conf` file.
- The security of file systems is enhanced by allowing the permission and ownership of certain files and directories to certain groups or people. This enhancement gives certain protection against the discovery of future vulnerabilities such as eliminating common system administrator errors. It also assigns root ownership of certain directories to prevent any intentional non-root user access. Security hardening capability promotes secure password choices by setting restricted default user attributes such as setting the default `umask` to `077` which allows only the files' owner to read all newly created files to protect confidentiality of the user data.
- The security hardening capability provides the configuration of network security options to protect against common network denial of service attacks. Also it disallows remote root logins and root FTP transfers to prevent hackers from sniffing root passwords on the network. Another advantage of disallowing remote root logins is that one has to `switchuser` to gain root authority, and this type of action is easily audited and monitored for any access violation.

7.3.3 IBM WebSphere Application Server Version 2.02

This application is provided on the AIX Bonus Pack shipped with AIX Version 4.3.3.

WebSphere Application Server Version 2.0, Standard Edition, takes you from simple Web publishing to powerful e-business applications. New features include:

- Web application deployment through dynamic content generation using database access and Java Server Pages (JSP).
- e-business solutions that grow with your business needs.
- Reliable deployment and management of server-side applications and components.
- An integrated tools environment for creating dynamic content and Java-based applications for deployment on all WebSphere compliant server platforms.
- An easy-to-use administrative interface.

WebSphere Application Server Version 2.0, Standard Edition, is available with National Language Support. Menus, messages, and publications have been translated into the following languages:

- Spanish
- Brazilian Portuguese
- Italian
- German
- French
- Traditional Chinese
- Simplified Chinese
- Japanese
- Korean
- English

WebSphere Application Server Version 2.0, Standard Edition, combines the control and portability of server-side business applications with the performance and manageability of Java technologies to offer a comprehensive Java-based Web application platform. Version 2.0 is integral in managing and integrating enterprise-wide applications while leveraging

open Java-based technologies and APIs. It enables powerful interaction with enterprise databases and transaction systems.

This Web application server, which deploys and manages Java applications, includes these features:

- Full Java support across multiple operating platforms.
- Quick and easy installation across popular Web servers.
- Easy-to-use servlet management and control.
- Full support for the latest Java APIs for servlets and session tracking.
- Support for dynamic Web content generation, including use of Java Server Pages (JSP), servlet applications, and Java beans.
- A function-rich IBM HTTP server (based on technology from the Apache HTTP server) with additional SSL-based security and performance features.
- Tivoli-ready modules that can be managed by Tivoli-based tools.
- Enhanced administration with an improved user interface for application deployment and server administration.
- XML Document Structure Services that enable e-business applications that interchange data via XML to parse, generate, manipulate, and validate XML and XSL-based content.

This product includes software developed by the Apache Group for use in the Apache HTTP server project. For details, visit their Web site at:

<http://www.apache.org>

WebSphere Application Server Version 2.0, Standard Edition, with its open and flexible Web application run-time environment, encompasses Java-based servers for Java-based applications and components. Version 2.0 includes:

- Simple installation
- Better performance than other methods of extending server function, running as an in-process plug-in with Apache, Netscape, Microsoft Internet Information Server, Lotus Domino Go Webserver, and the HTTP server
- Full support for the latest session-tracking APIs
- A graphical interface for servlet management to:
 - Set options for loading local and remote servlets
 - Set initialization parameters

- Specify servlet aliases
- Create servlet chains and filters
- Log servlet messages
- A servlet sandbox feature for greater security when running remotely loaded servlets
- Support for dynamic page content for defining Java server pages containing embedded Java code and servlet tags
- NCSA tags and Java beans
- Servlet chaining
- MIME-based data filtering
- Loading from Jar files
- Dynamic reloading of modified servlets
- Recursive servlet and NCSA tag processing
- Aliasing

WebSphere Application Server Version 2.0, Standard Edition, plugs into the following servers:

- Apache Server Version 1.3.2, and later, for Windows NT, Sun Solaris, and AIX
- Netscape Enterprise Server Version 3.0, and later, for Windows NT, Sun Solaris, and AIX
- Netscape FastTrack Server Version 3.0, and later, for Windows NT, Sun Solaris, and AIX
- Microsoft Internet Information Server Version 3.x and 4.0
- Lotus Domino Go Webserver Release 4.6.2.5, and later, for Windows NT, Sun Solaris, and AIX
- HTTP Server V1.3.3 for Windows NT, Sun Solaris, and AIX

To make development and testing of applications easier, you can download a development copy of the Java servlet engine from the Web site:

<http://www.software.ibm.com/webservers>

This copy is for use in the development and testing, but not deployment, of Web applications and no support is provided. Users are required to agree to these terms before completing the download. IBM reserves the right to withdraw this offer at any time without further notice.

WebSphere Application Server V2.0, Advanced Edition, supports the Enterprise JavaBean standard and CORBA. Customers who want to migrate their CORBA applications from WebSphere Application Server Version 1.0 or Version 1.1 should consider the Advanced Edition.

7.3.4 WebSphere Performance Pack Version 2.0

WebSphere Performance Pack Version 2.0 enables greatly enhanced Web site availability and infrastructure performance with its integration of key standard TCP/IP technologies:

- Web Traffic Express — Robust proxy caching and filtering to significantly improve response times and costs of accessing Web servers
- SecureWay Network Dispatcher — Highly effective rules-based TCP/IP traffic routing and load balancing for improved performance and continuous availability
- AFS — Location-independent Web content file creation and administration, providing scalable and high availability using non-disruptive content replication and universal Web access while supporting online, remote management

7.3.4.1 Web Traffic Express

The caching component of Performance Pack Version 2.0, Web Traffic Express, is improved and made more user-friendly through the following:

- Transparent Proxy

Web Traffic Express can be configured to work with a properly configured network router to intercept all HTTP requests on port 80. The request will be satisfied from the proxy's cache or forwarded to the original destination. This proxy transparency is available for HTTP requests on the AIX platform.

- Remote Cache Access

Performance is improved by having multiple instances of Web Traffic Express located near each other in network terms (on the same LAN segment or installed in an RS/6000 SP) sharing the same cached content. This results in more efficient use of cache space by eliminating redundant copies of the same data for multiple servers and increasing the cache hit rate because the shared cache will be bigger than each server's individual cache.

- HTTP/1.1 Compliance

With the full implementation of HTTP/1.1, Web Traffic Express identifies itself as an HTTP/1.1 server. Persistent connections are supported from

the client to Web Traffic Express and from Web Traffic Express to the origin server.

- **SNMP Subagent and MIB Support**

Web Traffic Express provides a Simple Network Management Protocol (SNMP) subagent and Management Information Base (MIB). The SNMP subagent lets you use any SNMP-capable network management system, such as TME 10 NetView or TME 10 Distributed Monitoring, to monitor your proxy server's health, throughput, and activity. The MIB data describes how the proxy server is being managed, reflects current and recent server status, and provides server statistics.

- **Customization Exits**

Ease of use is improved by providing additional hooks for the plug-in programmer to maintain state information, and HTTP parsing is now done by Web Traffic Express instead of the programmer writing the plug-in. Allowing multiple concurrent plug-in operations and providing more exit points provide a more flexible environment for plug-in customization.

- **Variant Caching**

When plug-ins are used to alter the original data, the resulting document can be cached as well as the original document. Translation from one language to another and from color to gray scale are examples of CPU-intensive tasks that would be beneficial to cache.

- **Performance Enhancements**

The cache architecture is restructured, providing more efficient mapping of URLs onto the cache filesystem. This provides for quicker retrieval of cached objects, more efficient use of disk space within the cache, and quicker cache garbage collection. The cache also uses write-behind techniques for greater throughput. Additionally, Web Traffic Express now includes caching of Domain Name Server (DNS) lookup results giving quicker response time and further reducing network load.

- **Enhanced Log Maintenance**

In order to effectively manage the space requirements for the four logs (error, access, proxy access, and cache access) generated by Web Traffic Express, two improved functions have been implemented: compression and purging. Compression can be specified for the logs by date. The logs are compressed and stored with no regard for storage constraints. Purging is done by date and size of each file. The maximum size for each log is set during configuration of Web Traffic Express. If, when the logs are purged by date, the maximum size for a log is still exceeded, daily logs are removed until the maximum size is no longer exceeded.

- Error Message Personalization

The error messages sent to the client's browser are kept in HTML pages. This allows an administrator to change or alter the message. This could be used to give additional information, additional instructions, or include things such as logos.

- FTP Proxy Enhancements

Three improvements made to FTP functions include:

- Files can be uploaded (stored on FTP server) as well as retrieved from the server.
- In the authentication process, the user's password is prompted for instead of being included in the URL. The password is transmitted using an HTTP authentication header that uses base64 encoding.
- A configuration directive allows the server administrator to specify whether FTP URLs will be treated as relative URLs or absolute URLs.

- Secure Request Filtering

Filtering for SSL secure requests is done by using Platform for Internet Content Selection (PICS) filtering on the non-secure home page of the URL and applying the receive or block decision to the secure request.

- Configuration Enhancements

Configuring the server using the browser, local or remote, is improved and more user friendly. The configuration screen is divided into three areas: navigation, workspace, and header. Navigation between different areas of the configuration process is always available in the navigation portion of the screen with a click of the mouse. The forms to be filled in are displayed in the workspace. The header portion contains a help button and a restart server button. The server can be restarted anytime with the server restart button. When the help button is clicked, another instance of the browser that does not interfere with the form being used displays the help information. The help information is divided into three types: field definition of the form being viewed, task oriented, and server general help.

- Proxy Auto Configuration Support

This function provides the capability to create a Proxy Auto Configuration (PAC) file that, when referenced by browsers, dynamically routes requests to available servers. If the client browser is set to Automatic Proxy Configuration and pointed to the PAC file, changes in proxies used by the browser are transparent to the client. This can be used by the server administrator to reroute requests when servers are down, to balance

workload, to send requests for specific URLs to specific proxies, or other reasons specific to a particular installation.

- **Cancel Control**

This function lets the administrator set a parameter that determines whether a content file will continue to be stored if the requestor connection is terminated. The parameter value represents a percentage of the file being received. If the amount of the file already received is equal to or greater than the value, the whole file will be received and stored.

7.3.4.2 SecureWay Network Dispatcher

The workload balancing component of Performance Pack Version 2.0, SecureWay Network Dispatcher, is improved and made more user friendly with the addition of the following:

- **Server Directed Affinity**

Users can direct requests from specified clients to specific servers. An API provided in Network Dispatcher is used to define the client-to-server relationship desired in the Network Dispatcher affinity tables. The user defines where server requests from the specified client are routed to. Subsequent to that, for as long as the user wants it to remain in effect, all requests from the client are routed to that server. The user has the capability of querying the relationships, adding relationships, or removing relationships from affinity tables.

- **Improved Logging and Statistics**

The logging and statistical data kept by Network Dispatcher is maintained in binary format. In addition, the command line and GUI are used to set parameters on how the binary data will be logged. A sample Java program is included that provides access to and manipulation of the log data.

- **Remote Configuration**

An authenticated link using the GUI as well as configuring at the Network Dispatcher server is now an option. This productivity enhancement facilitates the centralization of Network Dispatcher administration and control. The authenticated link between the client doing the remote configuration and the server uses a public key/private key pair for authentication. The private key must be distributed to the clients selected for remote configuration of the server.

- **Content-Based Forwarding**

An administrator can route requests based on URL content. The URL content of a request is matched to user-generated rules that determine which server, or cluster of servers, it will be routed to. If a user-generated

rule results in a choice of multiple servers, Network Dispatcher determines the best server for the job.

- Star Cluster

This catch-all cluster includes two functions: Configure multiple aliases on the Network Dispatcher machine and have them all use the same port/server configuration and make the Network Dispatcher machine the default route for some traffic and have it load balance all the Internet Protocol (IP) traffic for ports configured on the star cluster without configuring any aliases on the Network Dispatcher machine.

- Star Port

You can create a load balancing configuration for traffic to any port that has not been explicitly defined in your Network Dispatcher configuration. This could be used, for example, to load balance firewalls or discard requests for ports that have not been configured.

7.3.4.3 AFS

The file replication component of Performance Pack, AFS file sharing product, is enhanced with the inclusion of Version 3.4a for AIX and Sun Solaris and Version 3.5 for Windows NT. It provides a scalable, high-performance, reliable distributed file sharing system with robust security and a simple administrative model.

AFS lets you improve productivity and collaboration across remote departments. You can break down barriers created by different platforms, physical locations, and job functions to let workers share files securely and efficiently, regardless of where the information, or the work team, is located. AFS combines security, scalability, availability, and Web access for your enterprise to give users the confidence that business-critical information is not only accessible, but also accurate, up-to-date, and protected from unauthorized access by rigorous security mechanisms. And with AFS, you can equip your workforce with the distributed file sharing capabilities to maximize productivity, reduce time-to-market, reduce overall costs whether in a networked or Web-enabled systems environment, and maintain protection from unauthorized access. You can extend your business applications beyond the enterprise to conduct e-business. Customers and suppliers can interact directly via the Web with your business processes and information.

To avoid lost productivity when a temporary failure occurs within a system, AFS replicates important information on multiple machines. AFS will automatically locate a replica of the information and re-establish the link. It also enables system administrators to perform many system maintenance

and reconfiguration tasks, while information remains online and accessible. Users rarely have to go without access to the information they need.

A graphical administration tool, the AFS Control Center, extends global file sharing in computing environments and simplifies administration tasks. The AFS Control Center helps simplify system administration of AFS servers by letting administrators manage entire environments from a single Windows management station. With the AFS Control Center, administrators can now more effectively monitor file server utilization, transfer collections of files among servers, and perform load balancing to improve overall system performance.

The AFS Control Center running on a Windows NT desktop provides a GUI-based tool for a single point of administration. It lets systems administrators perform easy, remote administration of AFS servers.

The administrative functions include:

- Drag-and-drop functionality
- Fileset quotas/thresholds and alerts system
- Continuous monitoring of server status
- Automated file set replication
- Create/delete, lock/unlock, and move file sets
- Create and configure aggregates and file sets
- Create, manage, and monitor processes
- Add/delete principals and groups from administration lists

7.3.5 SecureWay Firewall for AIX V4

To secure your network for e-business, firewalling, controlling access to and from a network, is an important part of your network security. IBM SecureWay Firewall Version 4 is part of the IBM SecureWay Software offerings, bringing broad security solutions to protect your business assets. IBM SecureWay Firewall for AIX delivers new management and security functions for using the Internet.

Highlights of this product include:

- Supports all three standard firewall architectures
- Provides the Security Dynamics ACE Server with a two-user license at no additional charge
- Centrally manages and configures multiple firewalls

- Provides easy and secure administration through AIX, Windows 95, and Windows NT clients
- Eliminates the need for expensive leased lines by using the Internet as a virtual private network
- Reduces cost through predefined services for fast setup
- Disables unsafe applications to ensure a secure platform for the firewall
- Proactively scans the firewall and other hosts to find potential security exposures
- Provides seamless Internet access using standard client software

Table 30 shows features and benefits of this product.

Table 30. Features and Benefits of SecureWay Firewall

Feature	Benefit
Strong authentication with ACE Server	Eliminates the security exposure of clear text passwords
Java-based administration	<p>Securely provides support for AIX, Windows 95, and Windows NT clients</p> <p>Provides for different levels of administrator authority and logs activities by ID</p> <p>Provides integrated product documentation and context-sensitive help</p> <p>Uses objects to easily represent subnets and servers in the network</p>
Virtual private networking	<p>Allows users to securely exchange information with other companies</p> <p>Reduces the need for expensive leased lines by using the Internet</p>
Enterprise Firewall Manager	<p>Centrally manages multiple firewalls from a single point</p> <p>Reduces the cost and complexity of managing multiple firewalls</p> <p>Allows a company to set one security policy for all Internet access</p>

Feature	Benefit
Hardening	<p>Provides a secure platform for the firewall</p> <p>Reduces configuration errors that can lead to security exposures</p> <p>Eliminates security holes and unneeded services</p>
Network Security Auditor	<p>Increases security by proactively scanning the network for security exposures</p> <p>Allows the administrator to easily find configuration errors</p> <p>Generates easy-to-read HTML-formatted reports Isolation of subnets</p> <p>Easily keeps multiple subnets within the network securely separated from each other</p> <p>Provides the ability to set policy on how various subnets can communicate</p> <p>Allows for a Demilitarized Zone (DMZ) to keep Internet users from entering the secure network</p>
Reporting/logging	<p>Provides the data to allow you to manage your Internet connection</p> <p>Identifies usage patterns</p>
SOCKS Version 5	<p>Provides users with seamless Internet access</p> <p>Is supported through standard client software</p>
Three firewalls in one	<p>Provides support for all three standard firewall architectures — expert filters, application proxy, and circuit-level gateway</p>
Alerting	<p>Provides automatic and immediate notification of security breaches</p>
Expert filters	<p>Controls traffic by time-of-day, address, port number, direction, and protocol</p> <p>Easily configures the firewall to disallow inappropriate bandwidth-consuming traffic</p> <p>Controls the flow of traffic at a granular level</p>

Feature	Benefit
SecureMAIL	A reinvented foundation for the mail relay Provides standards-based SMTP relay
National language support	Is translated into French, Spanish, Italian, Brazilian Portuguese, traditional Chinese, simplified Chinese, Japanese, and Korean Allows administrators to use their native languages Makes all languages available through each configuration client

7.3.6 AIX Fast Connect for Windows

AIX Fast Connect for Windows Version 2.1.0 enables AIX to be a part of the Microsoft Network Neighborhood. Windows clients have AIX file and print capabilities without the need for additional software on the clients. It uses Microsoft Common Internet File System (CIFS) and Server Message Block (SMB) protocols over TCP/IP.

Windows clients can access:

- AIX Journaled File System (JFS)
- CD File System (CDFS)
- Network File System (NFS) mounted subsystems
- AIX printing services

AIX Fast Connect for Windows includes:

- Support for Windows for Workgroups (WFW), Windows 95, Windows 98, and Windows NT 4.0 Workstation clients.
- CIFS client resource browsing that enables users to view available shared AIX files and printing services.
- Support for CIFS long file names.
- NetBIOS Request for Comments (RFC) 1001/1002 allows programs written for a NetBIOS environment to run over TCP/IP; NetBIOS is for use with the Fast Connect features only and is not available for other applications on AIX.
- Use of the TCP/IP Domain Name System (DNS) to resolve NetBIOS machine names.

- Microsoft Windows Name Server (WINS) support.
- Opportunistic locking.
- Unicode use, file, and printer name support.
- AIX authentication/authorization with encrypted passwords using 56-bit encryption.
- Interaction with Windows NT Server Version 4.0 to:
 - Provide PC user authentication and authorization
 - Work with the Windows NT Server domain master browser to find and publish shared resources across TCP/IP subnets
 - Work with WINS to resolve NetBIOS machine names

AIX Fast Connect for Windows also assumes the presence of at least one Windows NT server in the network environment to provide Microsoft Master Browser support.

Chapter 8. Linux on RS/6000

The Model B50 supports AIX Version 4.3.2 (by special request), 4.3.3, and also Linux. Please note that while IBM provides Model B50 hardware and AIX software support, it does not provide software support for Linux. Keep in mind that when people refer to Linux, they usually mean Linux for Intel platforms. There are ports for Linux that run on PowerPC, and that will be the basis for this discussion.

8.1 What is Linux?

In 1991, Linus Torvalds, a student at the University of Helsinki in Finland, developed, as a hobby, an operating system he called Linux. Torvalds' hobby is now an operating system that has developed with assistance from programmers around the world. Today, Linux is a complete multiuser UNIX operating system capable of multitasking and other features common to other operating systems, such as Windows NT. At first, it was developed on the Intel x86 architecture platform. But from its early days, Linux also supported many other architectures, such as PowerPC, SPARC, and Alpha. Compared to operating systems used today, Linux is very modest when it comes to hardware requirements. Linux is sometimes thought of as an integrated system and seems to be a whole package like AIX or Solaris. However, Linux is just the name of a kernel and there are many other parts, such as libraries, commands, and applications, that build on this kernel to make it more useful. Many of these parts are developed or distributed under the license of GNU Public License (GPL).

Initially, Linux was supported only by communities of developers or users themselves, mainly through the Internet. Now, many vendors offer services for hardware, software, and support.

8.2 LinuxPPC for RS/6000

Native porting of the Linux kernel to the PowerPC processor was developed by recognized experts in the Linux community and there is now a packaged Linux 2.2 Kernel for RS/6000. Basic hardware and IO (SCSI, Ethernet, 2D graphics) support is available for LinuxPPC that runs on the RS/6000 platform. Initial RS/6000 support that is provided by Yellow Dog Linux will be available for the 7043-150, 7025-F50, and 7046-B50. Distribution and support is provided through Web downloads, FAQs, and discussion forums. The following sections provide details on the support, features, and installation of Linux on RS/6000-enabled platforms.

8.3 Support for LinuxPPC on RS/6000 Platform

Support for LinuxPPC on RS/6000 platform is provided by the following short list of vendors or organizations:

- Terra Soft Solutions, Inc. provides installation support for Yellow Dog Linux Champion Server running on an RS/6000 7046-B50, 7025-F50, or 43P-150. For more information, refer to their Web site at:

<http://www.yellowdoglinux.com/>

- Linuxcare Inc. announced that they will provide fee-based support for LinuxPPC on select RS/6000 platforms 24-hours a day, 7-days a week. Linuxcare focuses on enterprise markets by providing commercial Linux support, consulting, and education services:
 - Linux Technical Support and Consulting provides call center and on-site support as well as the Linuxcare Knowledgebase, a collection of Linux support information.
 - Linux University teaches users how to get the most out of Linux.
 - Linuxcare Labs provide porting assistance, optimization, benchmarking, and certification of Linux software and hardware.
 - Open Source Development including CareTraker, a support incident management system.

For more details, see the following Web site at:

<http://www.linuxcare.com/>

- General information about LinuxPPC is usually applicable for LinuxPPC on the RS/6000 platform. You can view information about Linux on the PowerPC platform at the following Web site:

<http://www.linuxppc.org/> or <http://www.linuxppc.com/>

- RedHat, Inc. is a one of the major packagers and solution vendors for Linux. You can obtain general information about Linux at the following location:

<http://www.redhat.com/>

- IBM support is currently through Linux communities. Hardware support from IBM is the same as all RS/6000 products. To see the latest information about IBM support for Linux, you can access the following IBM Web site at:

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

8.4 Applications on LinuxPPC

There are many commercial and noncommercial applications for Linux. However, there are not as many commercial applications for LinuxPPC. You can obtain several applications to serve your needs by using *open source* noncommercial application products. This section describes some features and applications of the typical Linux package that provides Internet services.

8.4.1 Basic Features for Internetworking on Linux 2.2

The following basic features are provided in the Yellow Dog Linux Champion Server 1.1 that is available for the Model B50:

- POP3/IMAP Server
 - Ability to append messages to a remote folder
 - Ability to set standard and user-defined message status flags
 - Support for simultaneous update and update discovery in shared folders
 - New mail notification
- Sendmail V8.9
Sendmail Version 8.9.3 with an anti-spamming feature
- FTP Server
ProFTPD-improved easy configuration and enhanced security (included in Yellow Dog Linux V1.1)
- NIS/NFS
NIS/NFS compatible implementation
- IPv4
IPv4 supported under Linux V2.2 kernel and porting IPv6 for Linux kernel V2.2 under development
- Java Support
The Java Development Kit for LinuxPPC is available at <http://www.blackdown.org/java-linux/>. The latest version of the Blackdown Java-Linux port of JDK 1.1.6 to Linux and JDK 1.2 porting is currently advancing
- DNS
BIND 8.2.6 (included in Yellow Dog Linux 1.1)

8.4.2 The Apache HTTP Server

Currently, Apache is the most popular HTTP server in ISPs and other Web-hosting companies. Its origin is a popular HTTP server named NCSA httpd 1.3 and was rewritten by the Apache Group. Apache is widely recognized as a substantially faster and more stable server. The Apache HTTP server is an HTTP/1.1 compliant server and easy to configure and extend with third-party modules. Apache users can easily develop extension functions through the Apache module API modules. The Apache server is a Open Source product and full source code is distributed under license of GPL. Apache Group is now reformed to the Apache Software Foundation and is still developing Apache and related software such as Apache JServ JAVA servlet support version HTTP server.

The Apache HTTP server has following features:

- DBM databases for authentication
- Customized responses to error and problems
- Multiple Directory Index directives
- Unlimited numbers of Alias and Redirect directives
- Content negotiation
- Multi-homed servers

The latest version 1.3.9 the Apache has some new features as follows:

- Dynamic Shared Object supporting runtime module loading
- Internal implementation improvement about performance
- Apache Autoconf-style Interface enables easy configuration
- HTTP/1.1-style virtual hosts setting improvement
- Extensions of module function and module API

For more information, visit: <http://www.apache.org/> and <http://java.apache.org>

8.4.3 Squid - Firewall/Web Cache Server

The Apache HTTP Server Version 1.1 and above can also act as a caching and proxy server.

Squid is widely used as a proxy and caching server and includes support of HTTP, FTP, and DNS proxy. Local users using a Web client can access Web pages or FTP sites on the Internet through the Squid server. The Squid

server also has the ability to cache the contents through proxying. Squid can suppress the bandwidth of their networks.

The Squid server supports the following features:

- Proxying and caching of HTTP, FTP, and other URLs
- Proxying for SSL
- Cache hierarchies
- ICP, HTCP, CARRP, and Cache Digests compliance
- Transparent caching
- Cisco's WCCP support
- Extensive access controls
- HTTP server acceleration
- Management and observation through the network using SNMP

8.4.4 SAMBA - SMB/CIFS Compatible File/Printer Server

SAMBA is a program package that supports SMB/CIFS file sharing protocols. Originally, SMB and CIFS were used on DOS/Windows platforms, but SAMBA now enables UNIX-based platforms to access these machines. SAMBA also offers printer sharing function based on SMB/CIFS. The SAMBA packages has client capabilities, so you can use resources of Windows platform machine on your network. The SAMBA package offers these functions:

- File and Print services
You can access to/from Linux to/from Windows file systems and printer resources seamlessly.
- Authentication and Authorization
Logon capability to Windows NT Domain as a client and partial support of NT Domain Controller compatible implementation.
- Name Resolution
Windows Internet Name Service (WINS) compatible NetBIOS name resolution over TCP/IP.
- Service Announcement (Browsing)
SAMBA can act as a Local Master Browser then keep and serve a list of available services on the local network to other SMB clients.
- GUI base setup and management (SWAT)

Samba Web Administration Tool (SWAT) allows you to configure and manage easily SAMBA through network using your favorite Web client.

The latest SAMBA version is 2.0.5a. For more information about SMABA, access <http://samba.org/> and follow the links.

8.5 Installing LinuxPPC on the Model B50

This section describes how to install Yellow Dog Linux on RS/6000.

We installed Yellow Dog Linux on a RS/6000 F50 two-way SMP Model with a simple configuration. The installation was later performed on a Model 150 and a model B50. The first installation, and the source for this document, was done on a machine with the specifications provided in Table 31

Table 31. Test Machine Specifications

Test Machine Specifications	
Hardware	
CPU	PowerPC 604e 332 MHz two-way SMP
Memory	512 MB (256 MB recognized)
Network interface	Integrated Ethernet 10/100 Mbps
Disk	9.1 GB internal SCSI Disk Drive
Graphics	GXT130P
Software	
Yellow Dog Linux Version 1.1	

8.5.1 Installation Considerations

Yellow Dog Linux is available at <http://www.yellowdoglinux.com/> (Terra Soft Solution Inc.). You can order or download it from this site.

Currently supported devices are restricted. You should check the available supported devices from the Web site before you start installation. At the time of writing, only the following adapters are supported:

- POWER GXT130P Graphics Accelerator (# 2830)
- POWER GXT120P Graphics Accelerator (# 2838)
- IBM 10/100 Mbps Ethernet PCI Adapter (# 2968)

At the time of writing, booting the installation CD-ROM was not available, so a boot floppy install disk was created.

After installation, you may see the message `Warning: /boot/System.map has an incorrect kernel version` during the boot process or after executing some command. If so, replace the `/boot/System.map` file to latest one to match your kernel version.

8.5.2 Installation Procedures

This installation procedure is similar to other RedHat-based Linux packages on Intel x86 platforms except for the booting and the tailoring of platform-specific devices.

8.5.2.1 Installation Preparation

Before you begin, you must create a set of disks.

You should get latest install and boot images (named `zImage.initr` and `zImage.chrp`) from:

`ftp://ftp.yellowdoglinux.com/pub/yellowdog/champion-1.1/ppc/install/RS6000/`

The file `zImage.initr` is for installing packages and the other, `zImage.chrp`, for the boot kernel image. Copy them each in DOS format using the name `zimage`. At the time of writing.

1. Format two DOS floppy disks.
2. Copy the `zImage.initrd` file to the floppy and change the name to `ZIMAGE`.
3. Copy the `zImage.chrp` file to the floppy and change the name to `ZIMAGE`.

If you wish to maintain a copy of AIX on your system, you can install on a unallocated scratch disk. You will have to modify the procedures shown below to properly allocate the space on the free disk.

The CD must be in place before you boot from the install floppy disk.

8.5.2.2 Booting from Install Floppy Image

Set the multiboot setting to floppy as first priority on SMS to allow you to boot with the boot floppy in your floppy drive. Hard disk boot was not available at the time of writing, but the development effort is ongoing.

1. Put your install floppy (originally named `zImage.initr`) in your RS/6000 diskette drive.
2. Power on and wait for the `E1F1` code in the status LEDs, then press **F8** to get to the Open Firmware command prompt.

3. When the prompt is displayed, enter the following two commands:

```
setenv boot-device floppy:\zimage
```

```
boot
```

4. The boot process should now start. The LED shows E14D -> LinuxPPC 2.2.10 and the Linux boot sequence is visible on the system display (console) as shown in Figure 24.




Figure 24. Yellow Dog Linux Welcome Screen

5. You may start the RedHat Linux installer in almost the same way as an x86 platform. Please see installation manual for more details.
6. Select a language, in this case, select **English** then press **Enter** (see Figure 25 on page 161).




Figure 25. Language Selection

7. Select keyboard type **PC**, not ADB, and select **us**, then press **Enter** on the panels shown in Figure 26 and Figure 27 on page 162.




Figure 26. Keyboard Type Selection - 1




Figure 27. Keyboard Type Selection - 2

- 8. You can install from local CDROM or through the network (only integrated Ethernet). We recommend selecting **Local CDROM** (see Figure 28). After this selection, a panel shown in Figure 29 on page 163 will appear.

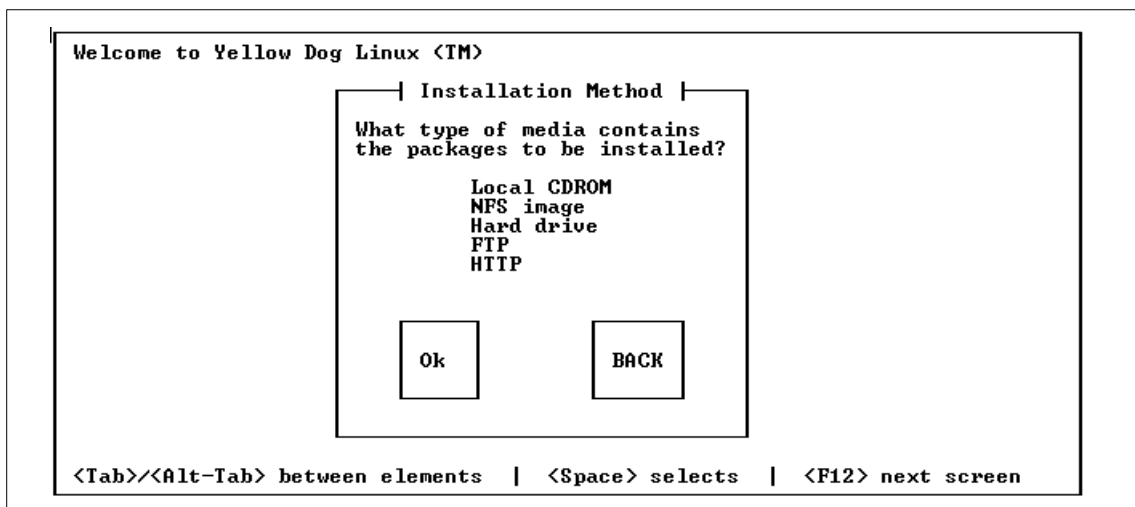


Figure 28. Installation Method - 1

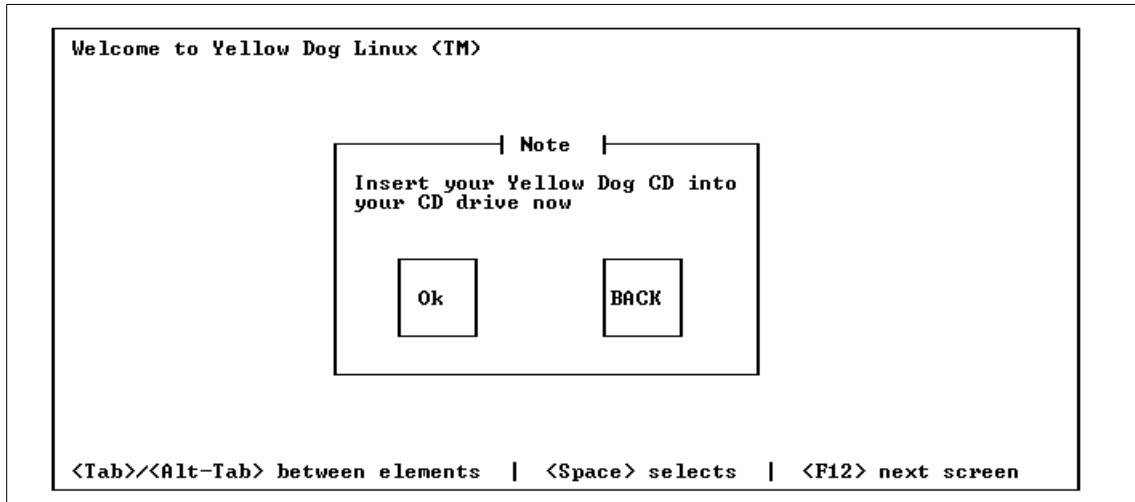


Figure 29. Installation Method - 2

9. During a test installation, the error messages in Figure 30 were shown, but the installation continued without interruption.

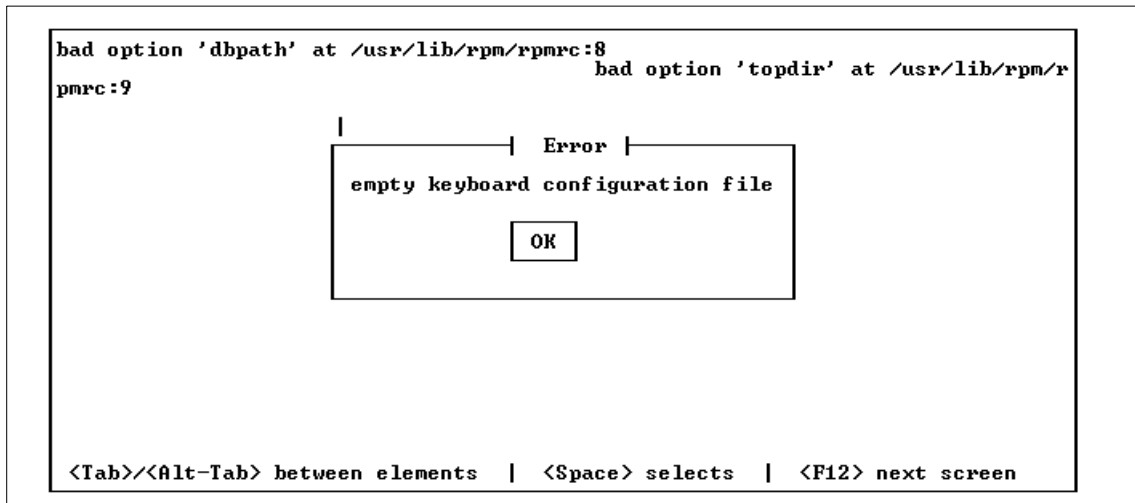


Figure 30. Empty Keyboard Configuration File Error Message

10. Select Install or Upgrade. We selected **Install** as shown in Figure 31.




Figure 31. Install/Upgrade Selection

11. You can see the disk drives on the display are those installed on your RS/6000. Select the drives that you want to install to. In this case, **/dev/hda** was selected. You must select **Edit** to partition the selected drive(s) (Figure 32 is the topic of the next discussion).




Figure 32. Disk Partitioning

12. You have two ways to partition disk(s), but using **fdisk** is more direct than using Disk Druid. When you finished partitioning then you can select **Done**.

- Press the **p** key to show the current partition settings.
 - Press **d** to select and delete a partition.
 - Press **n** to create a new partition.
 - Create the first primary partition. Press **n** then **p**, enter block numbers for first then last 1 and 20 (20 MB). This is for PReP boot.
 - Create the second primary partition. Press **n** then **p**, enter 21 and +1000MB as the value for the requirement. This will be the Linux / partition. We entered 21 and 2000 (about 2 GB).
 - Create the third primary partition. Press **n** then **p** to create appropriate system swap size for your RS/6000. We entered 2001 and +1024MB (1 GB for swap).
 - Set the partition type for each partition. Press **t** to select the type:
 - The first partition is set to 41 (PPC PReP Boot)
 - The second partition is set to 83 (Linux Native)
 - The third partition is set to 82 (Linux Swap)
 - You can create additional primary/extended partitions for other purposes.
1. When you are finished partitioning, check the current settings and then press **w** to write current settings and **q** to quit.

2. Here is an example procedure of how to make partitions:

```
Screen Setup filesystems - fdisk
```

```
This is the fdisk program for partitioning your drive. It is running
on /dev/sda.
```

```
The number of cylinders for this disk is set to 8678.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
```

```
Command (m for help): m
Command action
a  toggle a bootable flag
b  edit bsd disklabel
c  toggle the dos compatibility flag
d  delete a partition
l  list known partition types
```

```

m print this menu
n add a new partition
o create a new empty DOS partition table
p print the partition table
q quit without saving changes
s create a new empty Sun disklabel
t change a partition's system id
u change display/entry units
v verify the partition table
w write table to disk and exit
x extra functionality (experts only)

Command (m for help): p

Disk /tmp/sda: 64 heads, 32 sectors, 8678 cylinders
Units = cylinders of 2048 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
Command (m for help): n
  e extended
  p primary partition (1-4)
P
Partition number (1-4): 1
First cylinder (1-8676, default 1): <Enter>
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-8678, default 8678): 20

Command (m for help): n
  e extended
  p primary partition (1-4)
P
Partition number (1-4): 2
First cylinder (21-8676, default 21):
Using default value 21
Last cylinder or +size or +sizeM or +sizeK (21-8678, default 8678): +2000M

Command (m for help): n
  e extended
  p primary partition (1-4)
P
Partition number (1-4): 3
First cylinder (2022-8676, default 2022): <Enter>
Using default value 2022
Last cylinder or +size or +sizeM or +sizeK (2022-8678, default 8678): +1000M

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): L

 0 Empty          16 Hidden FAT16   61 SpeedStor     a6 OpenBSD
 1 FAT12          17 Hidden HPFS/NTF 63 GNU HURD or Sys a7 NeXTSTEP
 2 XENIX root     18 AST Windows swa 64 Novell Netware b7 BSDI fs
 3 XENIX usr      24 NEC DOS        65 Novell Netware b8 BSDI swap
 4 FAT16 <32M    3c PartitionMagic 70 DiskSecure Mult c1 DRDOS/sec (FAT-
 5 Extended      40 Venix 80286    75 PC/IX          c4 DRDOS/sec (FAT-
 6 FAT16         41 PPC PReP Boot  80 Old Minix      c6 DRDOS/sec (FAT-
 7 HPFS/NTFS     42 SFS            81 Minix / old Lin c7 Syrinx
 8 AIX           4d QNX4.x         82 Linux swap     db CP/M / CTOS / .
 9 AIX bootable  4e QNX4.x 2nd part 83 Linux          e1 DOS access
 a OS/2 Boot Manag 4f QNX4.x 3rd part 84 OS/2 hidden C:  e3 DOS R/O
 b Win95 FAT32   50 OnTrack DM     85 Linux extended e4 SpeedStor
 c Win95 FAT32 (LB 51 OnTrack DM6 Aux 86 NTFS volume set eb BeOS fs

```

```

e Win95 FAT16 (LB 52 CP/M      87 NTFS volume set f1 SpeedStor
f Win95 Ext'd (LB 53 OnTrack DM6 Aux 93 Amoeba      f4 SpeedStor
10 OPUS          54 OnTrackDM6      94 Amoeba BBT    f2 DOS secondary
11 Hidden FAT12  55 EZ-Drive      a0 IBM Thinkpad hi fe LANstep
12 Compaq diagnost 56 Golden Bow    a5 BSD/386      ff BBT
14 Hidden FAT16 <3 5c Priam Edisk

```

```

Command (m for help):
Hex code (type L to list codes): 41
Changed system type of partition 1 to 41 (PPC PReP Boot)

```

```

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 83

```

```

Command (m for help): t
Partition number (1-4): 3
Hex code (type L to list codes): 82
Changed system type of partition 3 to 82 (Linux swap)

```

```

Command (m for help): p

```

```

Disk /tmp/sda: 64 heads, 32 sectors, 8678 cylinders
Units = cylinders of 2048 * 512 bytes

```

Device	Boot	Start	End	Blocks	Id	System
/tmp/sda1		1	20	20464	41	PPC PReP Boot
/tmp/sda2		21	2021	2049024	83	Linux
/tmp/sda3		2022	3025	1025024	82	Linux swap

```

Command (m for help): w

```

- Next, edit the mount point for /dev/sda2 (for your / partition) and other mount point. We selected /dev/sda2 as / (see Figure 33 and Figure 34 on page 168).

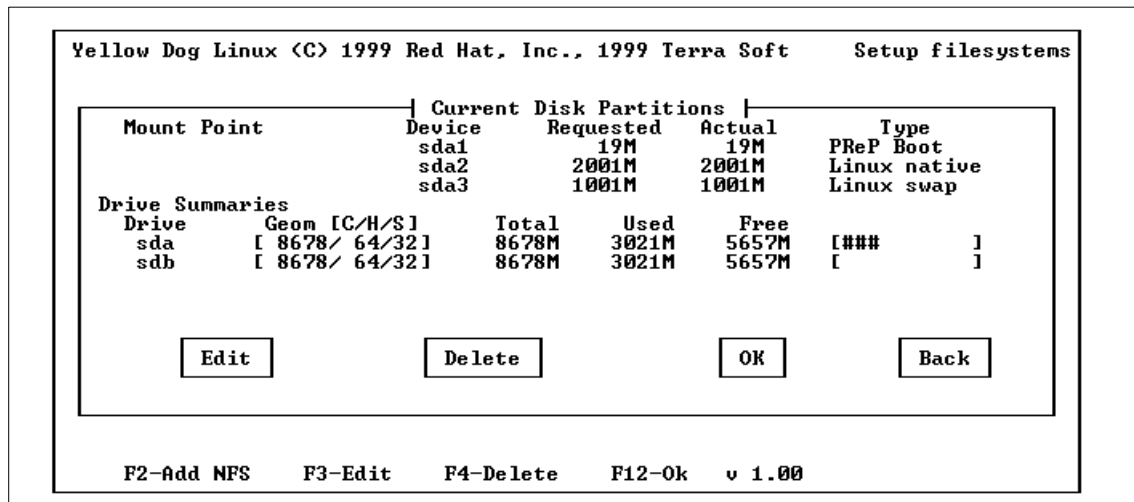


Figure 33. Filesystem Setup - 1




Figure 34. Filesystem Setup - 2

14. It is recommended that you reboot the machine at this point using **Ctrl+Alt+Del**, so the new partition settings made to the disk take effect.
15. When you restart your machine from the install floppy, recheck the partitioning to see that it has the proper settings (see Figure 35).




Figure 35. Swap Space Setting

16. Next, select the format option to check the disk parameter and partition formatting (see Figure 36).

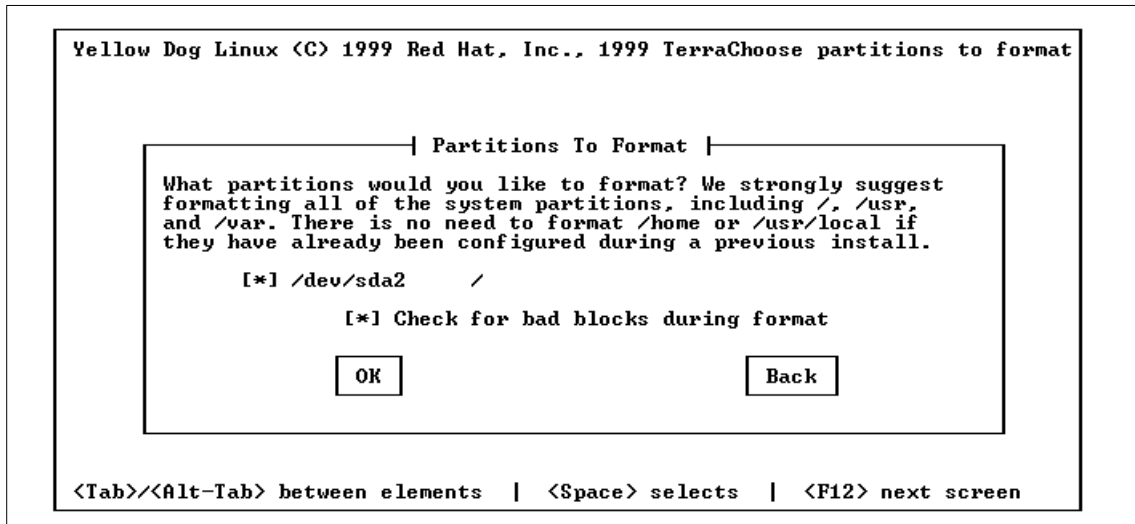


Figure 36. Formatting Partitions

17. Before formatting your / partition, you can select the packages you want to install (see Figure 37 on page 170). We chose entire packages by selecting **Everything**, then the formatting and installation of packages started. It required about 10 minutes to complete this task. You can see the installation process on your display.




Figure 37. Selecting Packages

18. You can view your installation log in the log file after the installation has completed (see Figure 38).

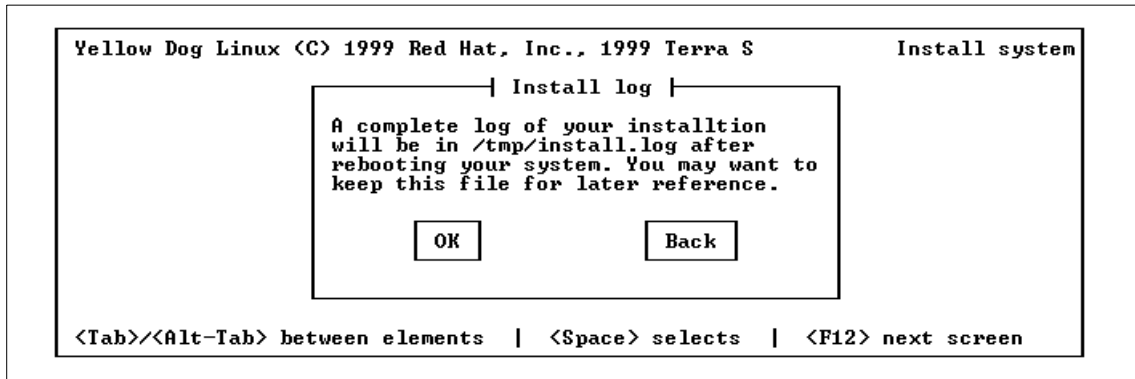


Figure 38. Installation Log

19. Select a mouse type. We chose a **Generic 3 Button Mouse (PS2)**. If you have a 3-button mouse, there is no need to check the Emulate 3 buttons? option (see Figure 39 and Figure 40).




Figure 39. Mouse Setting - 1




Figure 40. Mouse Setting - 2

20. Configuring your network is next. You need to know the IP address, netmask, hostname, if you have a default router, and nameserver. Set these values as shown in Figure 41, Figure 42, and Figure 43 on page 173.




Figure 41. Network Configuration - 1

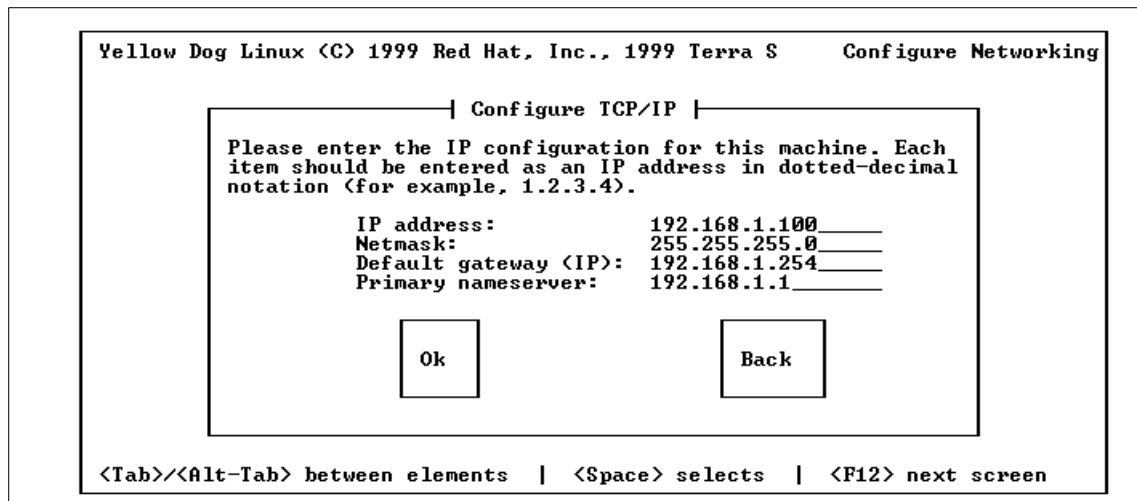


Figure 42. Network Configuration - 2




Figure 43. Network Configuration - 3

21. Next, configure a timezone (see Figure 44). We selected **US/Central** as our timezone.




Figure 44. Timezone Configuration

22. You can select services to run as daemons as shown in Figure 45 on page 174. By default, ftpd and telnetd are not configured.




Figure 45. Services Selection

23. If you have a printer, you can set printer parameters as shown in Figure 46.

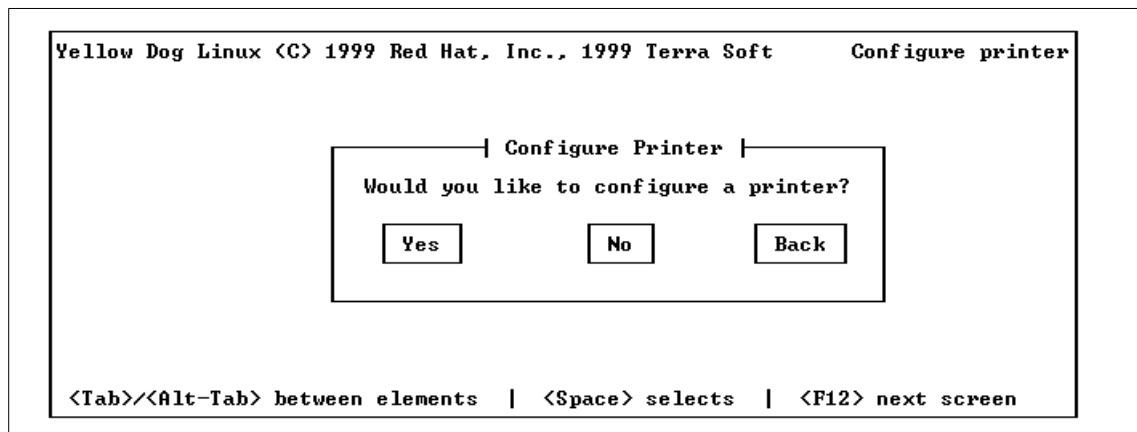


Figure 46. Printer Configuration

24. Enter and remember the root password (see Figure 47).




Figure 47. Password Setting

25. Choose **Authentication Configuration** (see Figure 48). If you need to integrate into a NIS domain, you can set it here.




Figure 48. Authentication Selection

26. At the time of writing, the quik 2.0 bootloader is available by FTP. If it is not installed, you may see the message shown in Figure 49.




Figure 49. Message About Bootloader

27. Select your monitor type and X window system configurations (see Figure 50 and Figure 51 on page 177). You can select the resolutions available on the target machine. When a test pattern is appropriately displayed on your screen, select the **Yes** button to complete.

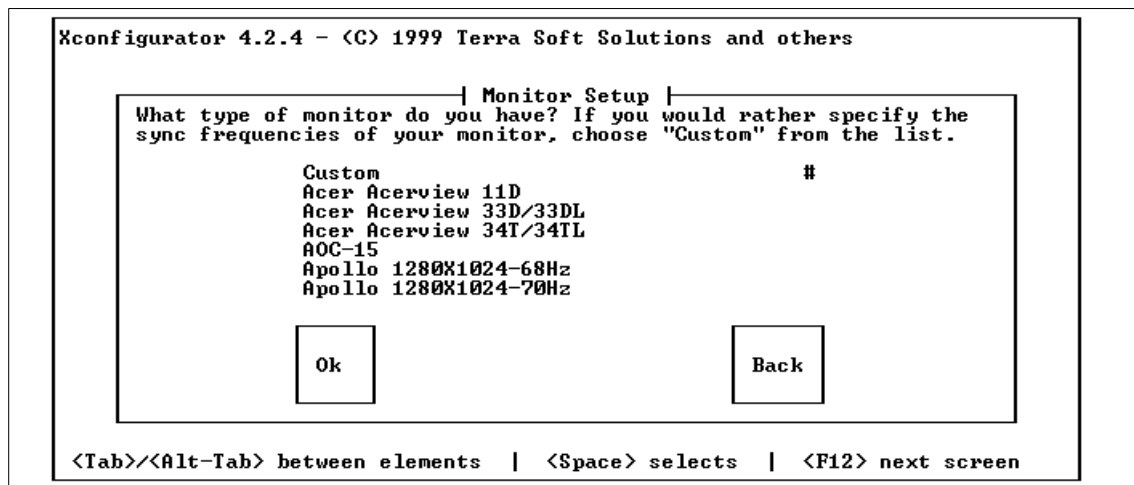


Figure 50. Monitor Selection/Setting

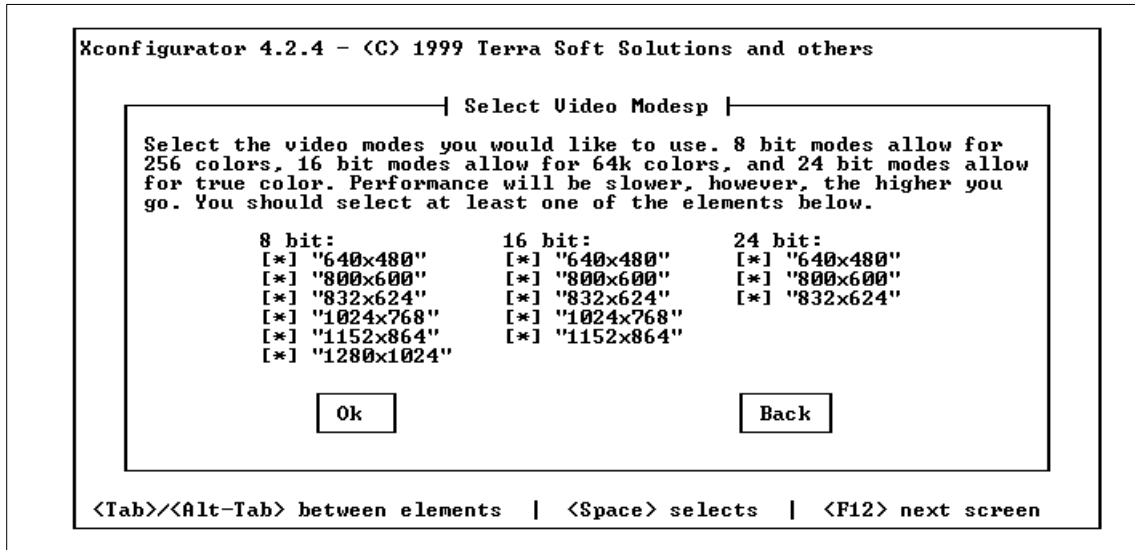


Figure 51. Resolution Selection for Display

28. You may see the message shown in Figure 52 when you finish your installation.




Figure 52. Linux Installation Complete Message

This should complete your installation. Do not forget to change the floppy to the boot floppy (originally named zImage.chrp). Then, reboot your system.

If you choose, install the quik 2.0 bootloader available at:

```
ftp://ftp.ppc.kernel.org/pub/linuxppc/quik/quik-2.0-0c.ppc.rpm
```

Once installed, this tool will set up your hard disk to boot your Linux system without the need for a boot floppy.

You can log on to Linux from your favorite desktop environment once the Linux system is restarted.

Here is a sample `dmesg` command message from this installation:

```
[root@ydl1 /proc]# dmesg
device tree used 37640 bytes
Total memory = 256MB; using 1024kB for hash table (at c0300000)
Linux version 2.3.10 (cort@chimaera.ppc.kernel.org) (gcc version egcs-2.91.62 19990224
(egcs-1.1.2 pre-release-2)) #5 Tue Jul 13 20:20:11 MDT 1999
Boot arguments:
Warning: no mac-io found
sio: keyboard irq level 255, type 255, active: remapping to level 1, type 2, active
sio: mouse irq level 255, type 255, active: remapping to level 12, type 2, active
PCI bus 0 controlled by pci at fef00000
PCI bus 64 controlled by pci at fee00000
PCI bus 128 controlled by pci at fed00000
RTAS Event Scan Rate: 1 (2999 jiffies)
OpenPIC Version 1.0 (4 CPUs and 16 IRQ sources) at f7fc0000
OpenPIC timer frequency is 4166666 Hz
time_init: decrementer frequency = 2490017940/60
Console: colour dummy device 80x25
Calibrating delay loop... 221.18 BogoMIPS
Memory: 254856k available (1496k kernel code, 5500k data, 144k init) [c0000000,d0000000]
Buffer-cache hash table entries: 16384 (order: 4, 65536 bytes)
Page-cache hash table entries: 65536 (order: 6, 262144 bytes)
POSIX conformance testing by UNIFIX
PCI: Probing PCI hardware
Warning: no ADB interface detected
Don't know how to access NVRAM with 4 addresses
Linux NET4.0 for Linux 2.3
Based upon Swansea University Computer Society NET3.039
NET4: Unix domain sockets 1.0 for Linux NET4.0.
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
TCP: Hash tables configured (established 262144 bind 262144)
Starting kswapd v 1.5
matroxfb: Matrox Mystique 220 (PCI) detected
matroxfb: 640x480x8bpp (virtual: 640x3270)
matroxfb: framebuffer at 0xDE000000, mapped to 0xd4806000, size 2097152
Console: switching to colour frame buffer device 80x30
fb0: MATROX VGA frame buffer device
Detected PS/2 Mouse Port.
PowerMac Z8530 serial driver version 1.01
pty: 256 Unix98 ptys configured
Macintosh ADB mouse driver installed.
Macintosh non-volatile memory driver v1.0
Sound initialization started
Sound initialization complete
RAM disk driver initialized: 16 RAM disks of 4096K size
loop: registered device at major 7
Uniform Multi-Platform E-IDE driver Revision: 6.19
Floppy drive(s): fd0 is 2.88M
```



```

FDC 0 is a National Semiconductor PC87306
md driver 0.36.6 MAX_MD_DEV=4, MAX_REAL=8
linear personality registered
raid0 personality registered
raid1 personality registered
raid5 personality registered
sym53c8xx: at PCI bus 0, device 12, function 0
sym53c8xx: setting PCI_COMMAND_IO PCI_COMMAND_MEMORY...
sym53c8xx: setting PCI_COMMAND_MASTER PCI_COMMAND_PARITY...(fix-up)
sym53c8xx: 53c825a detected
sym53c8xx: at PCI bus 2, device 11, function 0
sym53c8xx: setting PCI_COMMAND_IO PCI_COMMAND_MEMORY...
sym53c8xx: setting PCI_COMMAND_MASTER...(fix-up)
sym53c8xx: 53c825a detected
sym53c825a-0: rev=0x13, base=0x1ffae000, io_port=0x7fffc00, irq=17
sym53c825a-0: ID 7, Fast-10, Parity Checking
sym53c825a-0: on-chip RAM at 0x1ffaf000
sym53c825a-0: restart (scsi reset).
sym53c825a-0: Downloading SCSI SCRIPTS.
sym53c825a-1: rev=0x13, base=0x4fffe000, io_port=0x17fffc00, irq=26
sym53c825a-1: ID 7, Fast-10, Parity Checking
sym53c825a-1: on-chip RAM at 0x4ffff000
sym53c825a-1: restart (scsi reset).
sym53c825a-1: Downloading SCSI SCRIPTS.
scsi0 : sym53c8xx - version 1.3c
scsi1 : sym53c8xx - version 1.3c
scsi : 2 hosts.
  Vendor: IBM      Model: CDRM00203      !K Rev: 1_01
  Type:   CD-ROM          ANSI SCSI revision: 02
Detected scsi CD-ROM sr0 at scsi0, channel 0, id 4, lun 0
sym53c825a-0-<5,*>: FAST-10 SCSI 6.7 MB/s (150 ns, offset 8)
  Vendor: ARCHIVE  Model: IBM-STD224000N!D Rev: 726B
  Type:   Sequential-Access      ANSI SCSI revision: 02
Detected scsi tape st0 at scsi0, channel 0, id 5, lun 0
sym53c825a-0-<8,*>: FAST-10 WIDE SCSI 20.0 MB/s (100 ns, offset 8)
  Vendor: IBM      Model: DNES-309170W      Rev: SA40
  Type:   Direct-Access          ANSI SCSI revision: 03
Detected scsi disk sda at scsi0, channel 0, id 8, lun 0
  Vendor: IBMAUSTN Model: HSBP06AUSYSTEM    Rev: 0117
  Type:   Enclosure             ANSI SCSI revision: 03
sym53c825a-0-<8,0>: tagged command queue depth set to 8
sym53c825a-1-<8,*>: FAST-10 WIDE SCSI 20.0 MB/s (100 ns, offset 8)
  Vendor: IBM      Model: DNES-309170W      Rev: SA40
  Type:   Direct-Access          ANSI SCSI revision: 03
Detected scsi disk sdb at scsi1, channel 0, id 8, lun 0
  Vendor: IBMAUSTN Model: HSBP06AUSYSTEM    Rev: 0117
  Type:   Enclosure             ANSI SCSI revision: 03
sym53c825a-1-<8,0>: tagged command queue depth set to 8
scsi : detected 1 SCSI tape 1 SCSI cdrom 2 SCSI disks total.
sym53c825a-0-<4,*>: FAST-10 SCSI 10.0 MB/s (100 ns, offset 8)
Uniform CDROM driver Revision: 2.55
SCSI device sda: hdwr sector= 512 bytes. Sectors= 17774160 [8678 MB] [8.7 GB]
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 17774160 [8678 MB] [8.7 GB]
PPP: version 2.3.7 (demand dialling)
TCP compression code copyright 1989 Regents of the University of California
PPP line discipline registered.
pcnet32.c: PCI bios is present, checking for devices...
PCI Master Bit has not been set. Setting...
Found PCnet/PCI at 0x7fff800, irq 18.
eth0: PCnet/PCI II 79C970A at 0x7fff800, 00 04 ac 49 8a 22 assigned IRQ 18.
pcnet32.c:v1.21 31.3.99 tsbogend@alpha.franken.de
Partition check:
  sda: sda1 sda2 sda3 sda4

```

```
sdb:
VFS: Mounted root (ext2 filesystem) readonly.
Freeing unused kernel memory: 144k init 32k prep 4k pmac
Adding Swap: 1049596k swap-space (priority -1)
VFS: Disk change detected on device sr(11,0)
ISO 9660 Extensions: RRIP_1991A
Warning: defective CD-ROM (volume sequence number). Enabling "cruft" mount option.
```

Appendix A. Special Notices

This publication is intended to help sales specialists, DP managers, and hardware specialists evaluate and support the RS/6000 7046 Model B50 ISP Server. The information in this publication is not intended as the specification of any programming interfaces that are provided by the RS/6000 7046 Model B50. See the PUBLICATIONS section of the IBM Programming Announcement for the RS/6000 7046 Model B50 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have

been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

This document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

ADSTAR®	AIX®
AS/400®	DB2®
ESCON®	IBM®
Magstar®	Micro Channel®
Netfinity®	NetView
OS/2®	OS/390®
PowerPC Architecture	PowerPC 604®
QMF	RS/6000®
S/390®	SecureWay®
Service Director®	SP
System/390®	SystemXtra®
TURBOWAYS®	Ultrastar
WebSphere	

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere., The Power To Manage., Anything. Anywhere., TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

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

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

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

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

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

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Appendix B. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 IBM Redbooks Publications

For information on ordering these publications see “How to Get IBM Redbooks” on page 189.

- *IBM WebSphere Performance Pack Usage and Administration*, SG24-5233
- *Load-Balancing Internet Servers*, SG24-4993
- *RS/6000 43P 7043 Models 150 and 260 Handbook*, SG24-5144
- *RS/6000 Systems Handbook*, SG24-5120
- *The Technical Side of Being an Internet Service Provider*, SG24-2133
- *Implementing Netfinity Disk Subsystems: ServeRAID SCSI, Fibre Channel and SSA*, SG24-2098

B.2 IBM Redbooks Collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at <http://www.redbooks.ibm.com/> for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

B.3 Other Resources

These publications are also relevant as further information sources:

- *Customer Support Information*, SA23-2690
- *Diagnostics Information for Multiple Bus Systems*, SA38-0509
- *IBM RS/6000 Adapters, Devices, and Cables for Multiple Bus Systems*, SA38-0516
- *RS/6000 7046 B50 Setup Guide*, SA38-0562
- *RS/6000 7046 B50 User's Guide*, SA38-0563
- *RS/6000 7046 B50 Service Guide*, SA38-0564
- *System Unit Safety Information*, SA23-2652
- *IEEE 1275, IEEE Standards for Boot (Initialization Configuration) Firmware, Core Requirements and Practices*, ISBN 1-5593-7426-8
- *PowerPC Architecture: A Specification for a New Family of RISC Processors*, ISBN 1-5586-0316-6

B.4 Referenced Web Sites

These Internet sites are also relevant as further information sources:

- java.apache.org
- samba.org
- standards.ieee.org/index.html
- www.apache.org
- www.blackdown.org/java-linux
- www.elink.ibm.link.ibm.com/pbl/pbl
- www.ibm.com/linux
- www.isc.org/bind.html
- www.linuxcare.com/
- www.linuxppc.com
- www.linuxppc.org
- www.netlib.no/netlib/benchmark/performance.ps
- www.redbooks.ibm.com
- www.redhat.com

- www.rs6000.ibm.com/support/micro
- www.sendmail.org
- www.software.ibm.com/webserver
- www.specbench.org
- www.storage.ibm.com
- www.yellowdoglinux.com
- ftp.yellowdoglinux.com/pub/yellowdog/champion-1.1/ppc/install/RS6000/
- www.yellowdoglinux.com/products/preview_cs.shtml
- www.alphaworks.ibm.com/tech/spup
- www.rs6000.ibm.com/solutions/isp
- ftp.ppc.kernel.org/pub/linuxppc/quik/quik-2.0-0c.ppc.rpm

How to Get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** <http://www.redbooks.ibm.com/>

Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at <http://w3.ibm.com/> for redbook, residency, and workshop announcements.

Glossary

ABI	Application Binary Interface	CAM	Computer-Aided Manufacturing
AC	Alternating Current	CATIA	Computer-Graphics Aided Three-Dimensional Interactive Application
ADSM	ADSTAR Distributed Storage Manager		
ADSTAR	Advanced Storage and Retrieval	CD	Compact Disk
AIX	Advanced Interactive Executive	CD-ROM	Compact Disk Read Only Memory
ANSI	American National Standards Institute	CE	Customer Engineer
APAR	Authorized Program Analysis Report	CEC	Central Electronics Complex
API	Application Programming Interface	CHRP	Common Hardware Reference Platform
ASCI	Accelerated Strategic Computing Initiative	CLIO/S	Client Input/Output Sockets
ASCII	American National Standards Code for Information Interchange	CMOS	Complimentary Metal Oxide Semiconductor
ASP	Application Service Provider	COLD	Computer Output to Laser Disk
ATM	Asynchronous Transfer Mode	CPU	Central Processing Unit
BFF	Backup File Format	CRC	Cyclic Redundancy Check
BOS	Base Operating System	CSR	Customer Service Representative
BI	Business Intelligence	CSS	Communication Subsystems Support
BIST	Built-In Self-Test	CSU	Customer Set-Up
BLAS	Basic Linear Algebra Subprograms	CSU	Channel Service Unit
BOS	Base Operating System	CWS	Control Workstation
CAE	Computer-Aided Engineering	DAS	Dual Attach Station
CAD	Computer-Aided Design	DASD	Direct Access Storage Device (Disk)
		DAT	Digital Audio Tape
		DC	Direct Current
		DDC	Display Data Channel

DDS	Digital Data Storage	F/C	Feature Code
DE	Dual-Ended	FC-AL	Fibre Channel-Arbitrated Loop
DFS	Distributed File System	FCP	Fibre Channel Protocol
DIMM	Dual In-Line Memory Module	FDDI	Fiber Distributed Data Interface
DIP	Direct Insertion Probe	FDX	Full Duplex
DIVA	Digital Inquiry Voice Answer	FRU	Field Replaceable Unit
DLT	Digital Linear Tape	FTP	File Transfer Protocol
DMA	Direct Memory Access	F/W	Fast and Wide
DOS	Disk Operating System	GPFS	General Parallel File System
DRAM	Dynamic Random Access Memory	GUI	Graphical User Interface
DSU	Data Service Unit	HACMP	High Availability Cluster Multi Processing
DW	Data Warehouse	HACWS	High Availability Control Workstation
EC	Engineering Change	HDX	Half Duplex
ECC	Error Checking and Correction	HIPPI	High Performance Parallel Interface
EEPROM	Electrically Erasable Programmable Read Only Memory	HiPS	High Performance Switch
EIA	Electronics Industry Association	HiPS LC-8	Low-Cost Eight-Port High Performance Switch
EISA	Extended Industry Standard Architecture	HP	Hewlett-Packard
ELA	Error Log Analysis	HPF	High Performance FORTRAN
EMIF	ESCON Multiple Image Facility	HPSSDL	High Performance Supercomputer Systems Development Laboratory
EPOW	Environmental and Power Warning	HP-UX	Hewlett-Packard UNIX
ESCON	Enterprise Systems Connection (Architecture, IBM System/390)	HTTP	Hypertext Transfer Protocol
ESSL	Engineering and Scientific Subroutine Library	Hz	Hertz
ETML	Extract, Transformation, Movement and Loading	IA	Intel Architecture
		IBM	International Business Machines Corporation

ID	Identification	LAPI	Low-Level Application Programming Interface
IDE	Integrated Device Electronics	LED	Light Emitting Diode
IDS	Intelligent Decision Server	LFT	Low Function Terminal
IEEE	Institute of Electrical and Electronics Engineers	LP	Linear Programming
I²C	Inter Integrated-Circuit Communications	LPP	Licensed Program Product
I/O	Input/Output	LVM	Logical Volume Manager
IP	Internetwork Protocol (OSI)	MAP	Maintenance Analysis Procedure
IPL	Initial Program Load	MAU	Multiple Access Unit
IrDA	Infrared Data Association (which sets standards for infrared support including protocols for data interchange)	Mbps	Megabits Per Second
		MBps	Megabytes Per Second
		MCA	Micro Channel Architecture
		MCAD	Mechanical Computer-Aided Design
IRQ	Interrupt Request	MES	Miscellaneous Equipment Specification
ISA	Industry Standard Architecture	MIP	Mixed-Integer Programming
ISB	Intermediate Switch Board	MLR1	Multi-Channel Linear Recording 1
ISDN	Integrated-Services Digital Network	MMF	Multi-Mode Fibre
ISP	Internet Service Provider	MP	Multiprocessor
ISV	Independent Software Vendor	MP	Multi-Purpose
		MPC-3	Multimedia PC-3
ITSO	International Technical Support Organization	MPI	Message Passing Interface
JBOD	Just a Bunch of Disks	MPP	Massively Parallel Processing
JFS	Journalled File System	MPS	Mathematical Programming System
JTAG	Joint Test Action Group		
L1	Level 1	MTU	Maximum Transmission Unit
L2	Level 2		
LAN	Local Area Network	MVS	Multiple Virtual Storage (IBM System 370 and 390)
LANE	Local Area Network Emulation		

MX	Mezzanine Bus	POSIX	Portable Operating Interface for Computing Environments
NCP	Network Control Point		
NFS	Network File System	POST	Power-On Self-Test
NIM	Network Installation Manager	POWER	Performance Optimization with Enhanced RISC (Architecture)
NT-1	Network Terminator-1		
NTP	Network Time Protocol		
NUMA	Non-Uniform Memory Access	PPP	Point-to-Point Protocol
NVRAM	Non-Volatile Random Access Memory	PREP	PowerPC Reference Platform
OCS	Online Customer Support	PSSP	Parallel System Support Program
ODM	Object Data Manager	PTF	Program Temporary Fix
OLAP	Online Analytical Processing	PTPE	Performance Toolbox Parallel Extensions
OS/390	Operating System/390	PTX	Performance Toolbox
OSL	Optimization Subroutine Library	PV	Physical Volume
OSLp	Parallel Optimization Subroutine Library	PVC	Permanent Virtual Circuit
P2SC	Power2 Super Chip	QMF	Query Management Facility
PAP	Privileged Access Password	QP	Quadratic Programming
PBLAS	Parallel Basic Linear Algebra Subprograms	RAM	Random Access Memory
PCI	Peripheral Component Interconnect	RAN	Remote Asynchronous Node
PDU	Power Distribution Unit	RAS	Reliability, Availability, and Serviceability
PE	Parallel Environment	RAID	Redundant Array of Independent Disks
PEDB	Parallel Environment Debugging	RDBMS	Relational Database Management System
PID	Program Identification	RIPL	Remote Initial Program Load
PIOFS	Parallel Input Output File System	ROLTP	Relative Online Transaction Processing
POE	Parallel Operating Environment	RPA	RS/6000 Platform Architecture
POP	Power-On Password	RVSD	Recoverable Virtual Shared Disk

RTC	Real-Time Clock	SPS	SP Switch
SAN	Storage Area Network	SPS-8	Eight-Port SP Switch
SAS	Single Attach Station	SRC	System Resource Controller
SAR	Solutions Assurance Review	SSC	System Support Controller
ScaLAPACK	Scalable Linear Algebra Package	SSA	Serial Storage Architecture
SCO	Santa Cruz Operations	STP	Shielded Twisted Pair
SCSI	Small Computer System Interface	SUP	Software Update Protocol
SDR	System Data Repository	SVC	Switch Virtual Circuit
SDRAM	Synchronous Dynamic Random Access Memory	TCL	Tool Command Language
SDLC	Synchronous Data Link Control	TCP/IP	Transmission Control Protocol/Internet Protocol
SE	Single-Ended	TCQ	Tagged Command Queuing
SEPBU	Scalable Electrical Power Base Unit	TPC	Transaction Processing Council
SGI	Silicon Graphics Incorporated	UDB EEE	Universal Database and Enterprise Extended Edition
SLIP	Serial Line Internet Protocol	UP	Uniprocessor
SLR1	Single-Channel Linear Recording 1	USB	Universal Serial Bus
SMIT	System Management Interface Tool	UTP	Unshielded Twisted Pair
SMS	System Management Services	UUCP	UNIX-to-UNIX Communication Protocol
SMP	Symmetric Multiprocessing	VESA	Video Electronics Standards Association
SOI	Silicon-on-Insulator	VG	Volume Group
SP	Scalable POWERParallel	VHDCI	Very High Density Cable Interconnect
SP	Service Processor	VM	Virtual Machine (IBM System 370 and 390)
SPEC	Standard Performance Evaluation Corp.	VMM	Virtual Memory Manager
SPOT	Shared Product Object Tree	VPD	Vital Product Data

VSD	Virtual Shared Disk
VSM	Visual Systems Management
VSS	Versatile Storage Server
VT	Visualization Tool
WAN	Wide Area Network
WTE	Web Traffic Express
XTF	Extended Distance Feature

Index

Numerics

100BaseTx Ethernet 30
10BaseT Ethernet 30
18.2 GB Ultra SCSI 16-bit Enhanced Disk Drive 84
2104 LVD expandable storage plus 85
2118 VHDCI to 68-pin adapter 36
2494 PCI 3-Channel Ultra2 SCSI RAID adapter 36
2624 32X (max) SCSI-2 CD-ROM drive 84
2708 Eicon ISDN DIVA PRO adapter 30
2709 ARTIC960Hx 4-Port E1 cable 35
2710 ARTIC960Hx 4-Port E1 RJ45 cable 35
2741 SysKonnnect FDDI-LP adapter (SAS) 31
2742 SysKonnnect FDDI-LP adapter (DAS) 31
2743 SysKonnnect FDDI-UP adapter (SAS) 32
2830 GXT130P graphics adapter 38
2943 8-port asynchronous adapter 34
2944 128-port asynchronous controller 34
2947 ARTIC960HX 4-port selectable adapter 35
2962 2-port multiprotocol adapter 34
2963 TURBOWAYS 155 UTP ATM adapter 33
2968 PCI 10/100 Mbps Ethernet adapter 30
2988 TURBOWAYS 155 MMF ATM adapter 33
3112 9.1 GB Ultra SCSI 16-bit disk drive 15
3113 18.2 GB Ultra SCSI 16-bit disk drive 15
3627 IBM P76 color monitor, stealth black 22
3628 IBM P260 color monitor, stealth black 22
3629 IBM P76 color monitor, pearl white 22
3630 IBM P260 color monitor, pearl white 22
3752 service package 26
3753 Options Package 27
4150 128 MB SDRAM DIMM (Model B50) 21
4169 256 MB SDRAM DIMM (Model B50) 21
4959 IBM High-Speed Token-Ring PCI adapter 29
5005 software pre-install 22
6041 3-button mouse, pearl white 22
604e design 97
6205 PCI Dual-Channel Ultra2 SCSI adapter 36
6206 PCI Single-Ended Ultra SCSI adapter 35, 68,
71, 72
6207 PCI Differential Ultra SCSI adapter 36, 68
6225 PCI Advanced SerialRAID adapter 37
6235 32 MB SSA Fast-Write Cache option card 37
6559 additional universal mounting rails 25
6XX-MX bus 95
7133-D40 Serial Disk System, Advanced 87
7133-T40 Serial Disk System, Advanced 87

8137 16-port enhanced remote EIA-232 node 34
8138 16-port enhanced remote RS-422 node 34
8741 3-button mouse, stealth black 22
9.1 GB Ultra SCSI 16-Bit Disk Drive 83

A

adapter performance 39
 ATM 43
 disk I/O throughput 44
 Ethernet 42
 network and disk performance, combined 46
 networking throughput 39
adapter placement guidelines
 Model B50 26
adapters
 128-port async 34
 16-port enhanced remote EIA-232 node 34
 8-port async EIA-232/RS-422 34
 ATM 33
 basic configuration 4
 communication 29
 communications
 performance notes 39
 Ethernet 30
 FDDI 31
 ISDN 30
 multiple port 34
 PCI SCSI-2 single-ended Ultra SCSI 35
 storage
 performance notes 39
 token-ring 29
 WAN 34
AFS 147
AIX Fast Connect for Windows 151
AIX4.3.3 134
Andrew File System 147
ANSI standard
 X3.131-1986 59
 X3T9.2/375R Revision 10K 60
Apache 141, 156
Apache HTTP server 141
Apache JServ 156
Apple Inc., regarding CHRP 98
architecture
 ISA 94
 PCI 93
 PowerPC 604e 96

- SCSI-III 61
- asynchronous communication, SCSI 65
- ATM adapters 33
 - ATM adapter performance 43
- attributes
 - key 4

B

- bays
 - Model B50 15
- Berkeley Internet Name Domain 131
- BIND 131
- BIND 8.1.2 136
- boot floppy
 - Linux 159
- boot sequence
 - in general 112
 - Model B50 113
- bus
 - 6XX-MX 95
 - ISA 96
- bus architecture 93

C

- cables, SCSI (various types) 67, 72
- cabling the PCI differential Ultra SCSI adapter 75
- cfgmgr, command 49, 54, 104
- Channel Service Unit 3
- CIFS 134, 151
- clearances 26
- command
 - cfgmgr 54, 104
 - diag 104, 107
 - lsattr 54
 - lscfg 54
 - shutdown 107
- common command set, SCSI 60
- Common Internet File System 151
- common Internet file system 134
- communication adapters 29, 30
 - ATM 33
 - FDDI 31
 - ISDN 30
 - multiple port 34
 - token-ring 29
 - WAN 34
- Compaq, regarding PCI involvement 93
- comparison between SSA and SCSI architecture

- 82
- configuration
 - adapters 55
 - minimum 3
- configuration database 49
- connectors
 - external 16
 - internal 13

D

- Data Service Unit 3
- device configuration database 49
- device location codes 49
- device-to-device cables 76
- diag, command 104, 107
- diagnostics
 - considerations 101
 - online 104
 - running from a TTY 102
 - standalone 103
 - tasks and service aids 108
- differential Ultra SCSI adapter, PCI 36
- differential, SCSI 66
- Digital (DEC), regarding PCI involvement 93
- directory services 133
- disconnect/reconnect, SCSI 60
- disk systems, external
 - 7133-D40 Serial Disk System, Advanced 87
 - 7133-T40 Serial Disk System, Advanced 87
 - Expandable Storage Plus 85
- DNS 131
- Domain Name Service 131

E

- Ethernet adapter 30
 - Ethernet adapter performance 42
 - integrated 17
 - Model B50 17
- external SCSI and SSA storage devices 85

F

- F1 key, boot function 117
- F5 key, boot function 117
- F6 key, boot function 118
- F8 key, boot function 118
- fast SCSI 64
- fast/wide, regarding SCSI-2 65

Fast/Wide, SCSI 65
 fast-20 wide
 SCSI 65
 fast-20, SCSI 65
 Fast-Write Cache option card 37
 FDDI adapters 31
 feature code
 18.2 GB Ultra SCSI 16-bit disk drive 84
 2104 LVD expandable storage plus 85
 2118 VHDCI to 68-pin adapter 36
 2494 PCI 3-Channel Ultra2 SCSI RAID adapter
 36
 2708 Eicon ISDN DIVA PRO 30
 2709 ARTIC960Hx 4-Port T1 cable 35
 2710 ARTIC960Hx 4-Port E1 RJ45 cable 35
 2741 SysKonnect FDDI-LP adapter 31
 2742 SysKonnect FDDI-LP adapter (DAS) 31
 2743 SysKonnect FDDI-UP adapter 32
 2830 GXT130P graphics adapter 38
 2943 8-port asynchronous adapter 34
 2944 128-port asynchronous controller 34
 2947 ARTIC960HX 4-port selectable adapter
 35
 2962 2-port multiprotocol adapter 34
 2963 TURBOWAYS 155 UTP ATM adapter 33
 2968 PCI 10/100 Mbps Ethernet adapter 30
 2988 TURBOWAYS 155 MMF ATM adapter 33
 3112 9.1 GB Ultra SCSI 16-bit disk drive 15
 3113 18.2 GB Ultra SCSI 16-bit disk drive 15
 32X (max SCSI-2 CD-ROM drive 84
 3627 IBM P76 color monitor, stealth black 22
 3628 IBM P260 color monitor, stealth black 22
 3629 IBM P76 color monitor, pearl white 22
 3630 IBM P260 color monitor, pearl white 22
 3752 service package 26
 3753 Options Package 27
 4150 128 MB SDRAM DIMM (Model B50) 21
 4169 256 MB SDRAM DIMM (Model B50) 21
 4959 IBM High-Speed Token-Ring PCI adapter
 29
 5005 software pre-install 22
 6041 3-button mouse, pearl white 22
 6205 PCI Dual-Channel Ultra2 SCSI adapter
 36
 6206 PCI single-ended Ultra SCSI adapter 35,
 68, 71, 72
 6207 PCI differential Ultra SCSI adapter 36, 68
 6225 PCI Advanced SerialRAID adapter 37
 6235 32 MB SSA Fast-Write Cache option card
 37
 6559 additional universal mounting rails 25
 8137 16-port enhanced remote EIA-232 node
 34
 8138 16-port enhanced remote RS-422 node
 34
 8741 3-button mouse, stealth black 22
 9.1 GB Ultra SCSI 16-bit drive 83
 PCI Single-Ended Ultra SCSI adapter 35
 File Sharing Service 133
 firewall 148
 firmware
 checkpoints 127
 introduction 112
 recovery 127
 update (SMS) 124
 FTP Service 133

G
 general SCSI consideration 68
 Gigabit Ethernet adapter
 performance of 42
 GNU Public License (GPL) 153
 GPL 153
 graphics adapter
 Model B50 22
 GXT130P, graphics adapter 38

H
 hard disk spin up delay 123
 hardware
 overview 91

I
 I/O ports 16
 Model B50 16
 IBM, regarding PCI involvement 93
 IDE, industry-standard interface 99
 IEEE 1149.1 98
 IEEE 1275 99
 IMAP4 134
 installation information 25
 Installation Procedures for LinuxPPC 159
 installation specifications 26
 integrated Ethernet adapter
 Model B50 17
 Intel, regarding PCI involvement 93

- internal CD-ROM drive
 - Model B50 21
- internal disk drives
 - Model B50 21
- internal SCSI devices 83
- internal SSA devices 83, 85
- Internet Service Provider (ISP) 129
- Internet Services 130
- Internet Services management and administration 138
- Internet services monitor 139
- Internet Software Consortium 131
- IPSEC 135
- IPv6 135
- IrDA, protocol 99
- ISA
 - bus architecture 94
- ISC 131
- ISDN adapters 30
- ISP 129
 - server topology 2
- ISP software 136

J

- Java 1.1.8 135
- Java Server Pages (JSP) 140
- JSP 140

K

- keyboard 21

L

- LDAP 135
- libsrcsec.so.zip 137
- LINPACK DP description 9
- LINPACK SP description 9
- LINPACK TPP description 9
- Linus Torvalds 153
- Linux 153
 - boot floppy 159
- Linux 2.2 Kernel 153
- Linuxcare Inc. 154
- LinuxPPC 153
- LocalTalk, industry-standard interface 99
- lsattr, command 54
- lscfg, command 54

M

- Mail Service 131
- mail transfer agent 132
- memory subsystem
 - Model B50 14
- Microprocessor Common Hardware Reference Platform 98
- minimum configuration 3
- Model B50
 - adapter placement guidelines 26
 - boot sequence 113
 - communications adapter 23
 - firmware checkpoints 128
 - graphics adapter 22
 - integrated Ethernet adapter 17
 - internal CD-ROM drive 21
 - internal disk drives 21
 - key attributes 4
 - memory subsystems 14
 - optional features 20
 - overview 3
 - PCI storage adapters 22
 - performance 9
 - processor subsystem 12
 - reliability, availability, and serviceability 18
 - standard features 12
 - supported devices 23
- modems 24
- monitors 22
- Motorola Inc., regarding CHRP 98
- mounting information 25
- mouse 22
- multiple port adapters 34

N

- NCR, regarding PCI involvement 93
- NEBS 10
- Network Dispatcher 146
- network file system 134
- NFS 134
- NIS+ 135
- NUMA 95

O

- Online Diagnostics 104
 - concurrent mode 106
 - service mode 105
 - single-user mode 107

- open firmware 118, 126
- optional features
 - Model B50 20
- Options Library (Hard Copy) 27
- overview
 - adapter 29
 - hardware 91

P

- P260, monitor 22
- P76, monitor 22
- PC card, protocol 99
- PCI
 - bus architecture 93
 - features and benefits 94
 - Special Interest Group 93
 - specification revision 2.0 93
 - specification revision 2.1 93
- PCI differential Ultra SCSI adapter 36
- PCI Single-Ended Ultra SCSI Adapter 35
- PCI single-ended Ultra SCSI adapter 35
- PCI slots
 - Model B50 16
- PCI storage adapters
 - Model B50 22
- performance
 - adapters, communications and storage 39
 - See also adapter performance
- performance summary 9
- physical location codes 49
- planar 12, 13
- Plotters 24
- POP3 132, 134
- ports 16
- POST 114
- Post Office Protocol Version 3 132
- power-on password 121
- power-on self test 114
- PowerPC
 - introduction 91
- PowerPC 604e microprocessor 96
- PowerPC Platform Specification (PRep) 98
- Printer Sharing Service 134
- printers
 - Model B50 24
- privileged access password
 - SMS 121
- processor

- mini-support 96
- PowerPC 604e 96
- processor subsystem
 - Model B50 12
 - POWER3 96
 - PowerPC 604e 96
- publications
 - Model B50 17

Q

- QoS 136
- Quality of Service (QoS) 136

R

- RAS 18
- RedHat, Inc. 154
- reliability, availability, and serviceability
 - Model B50 18
- repeaters, SCSI 66
- requirement from ISP and ASP marketplace 10
- requisites.tar 137
- RFC 1725 134
- RFC 1730 134
- RIPL 123
- RPA
 - achievements 99
 - description of 98
 - goals of 98
 - history of 98
 - introduction 91, 98
- RS/6000 7046 Model B50
 - hardware design 91
- RS/6000 Platform Architecture (RPA) 98
- RSVP 136

S

- SAMBA 157
- SCSI
 - asynchronous 65
 - common command set 60
 - differential 66
 - disconnect/reconnect 60
 - overview 59
 - repeaters 66
 - specifications summary 67
 - summary of specifications 67
 - synchronous 65

- table of modes/rates 67
- tagged command queuing 60
- terminators 78
- Ultra 64
- SCSI-III 61
- SecureMAIL 151
- SecureWay Firewall for AIX 148
- SecureWay Network Dispatcher 146
- security hardening 139
- Sendmail 132
- Sendmail V8.9 134
- server message block 134
- Service Package 26
- Service Providers Utility Pack 136
- shutdown, command 107
- single-ended Ultra SCSI adapter, PCI 35
- single-user mode 107
- SMB 134
- SMS 116
 - ASCII main menu 123
 - F1 key 117
 - F5 key 117
 - F6 key 118
 - F8 key 118
 - firmware checkpoints 127
 - firmware recovery 127
 - firmware update 124
 - graphical main menu 118
 - power-on password 121
 - privileged access password 121
 - start-up 116, 117
- SMTP 132
- SOCKS5 136
- Software Pre-Install 22
- Solaris 135
- SPECfp_base_rate95 description 8
- SPECfp_base95 description 8
- SPECfp_rate95 description 8
- SPECfp95 description 8
- specifications
 - installation 25
- SPECint_base_rate95 description 8
- SPECint_base95 description 8
- SPECint_rate95 description 8
- SPECint95 description 8
- SPECweb96 description 8
- spup.base 137
- spup.html.en_US 137
- Squid 156

- SSA
 - fast-write option card 37
- Standalone Diagnostics 103
- standard features
 - Model B50 12
- standards
 - ANSI X3.131-1986 59
 - X3T9.2/375R Revision 10K 60
- storage boot devices 59
- storage, external
 - disk
 - 7133-D40 Serial Disk System, Advanced 87
 - 7133-T40 Serial Disk System, Advanced 87
- Sun Solaris 135
- supported devices
 - Model B50 23
- synchronous, SCSI 65
- system board
 - connections 13
- system management services 116
- system planar 12

T

- tagged command queuing, SCSI 60
- Terminals 24
- terminators
 - SCSI 78
- Terra Soft Solutions, Inc. 154
- Tivoli-ready 141
- token-ring adapter 29
- tolerances
 - installation 26
- topology
 - server 2

U

- Ultra SCSI 64
- Ultra SCSI adapter, PCI differential 36
- Ultra SCSI adapter, PCI single-ended 35
- Ultra SCSI Wide versus Ultra SCSI 65
- Ultra SCSI, defined 65
- Ultra2 SCSI 65
- USENET 3

V

- virtual private networking 149

VPN 135, 149

W

WAN adapters 34
Web Service 133
Web Traffic Express 143
Web-Based System Management 56
WebSphere 140
WebSphere Application Server 2.02 140
WebSphere Performance Pack 143
Wide SCSI 65
wide SCSI 65
workstation family 1
workstation offering 1

X

X-bus 96
XML Document Structure Services 141

Y

Year2000 17
Yellow Dog Linux 153

Z

ZIMAGE installation file 159
zImage.chrp 159
zImage.initr 159

IBM Redbooks Evaluation

RS/6000 7046 Model B50 Handbook
SG24-7046-00

Your feedback is very important to help us maintain the quality of IBM Redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com/>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

Customer **Business Partner** **Solution Developer** **IBM employee**
 None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other Redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-7046-00
Printed in the U.S.A.

RS/6000 7046 Model B50 Handbook

SG24-7046-00

