RS/6000 S-Series Enterprise Servers Handbook

Richard Cutler, Ron Barker, Mauro Minomizaki

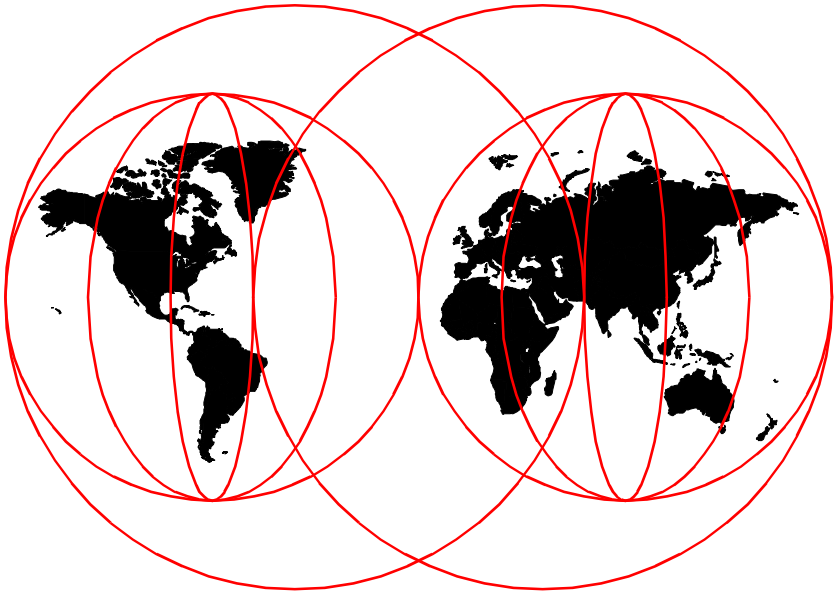

International Technical Support Organization

www.redbooks.ibm.com

SG24-5113-00

SG24-5113-00

International Technical Support Organization

RS/6000 S-Series Enterprise Servers Handbook

December 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special notices” on page 109.

First Edition (December 1999)

This edition applies to the RS/6000 Model S80 Enterprise Server and the RS/6000 Model S70 Advanced Enterprise Server for use with Version 4.3.3 of the AIX operating system, program number 5765-C34.

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.

Contents

Figures

Tables

Preface

This redbook covers the RS/6000 Model S80 and Model S70 Advanced Enterprise servers. It will help you understand the architecture of each machine and the similarities and differences between them. An overview of the optional features for each machine is also provided along with advice on how to use the PCRS6000 Configurator to produce a valid configuration.

This publication is suitable for professionals wishing to acquire a better understanding of RS/6000 S-Series Enterprise Servers including:

- Customers
- Sales and Marketing professionals
- Technical Support professionals
- Business Partners

This publication does not replace the latest RS/6000 marketing materials and tools. It is intended as an additional source of information that, together with existing sources, may be used to enhance your knowledge of IBM S-Series Enterprise Server products.

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.

Richard Cutler is an AIX and RS/6000 Technical Specialist at the ITSO, Austin Center. Before joining the ITSO, he worked in the RS/6000 Technical Center in the UK where he assisted customers and independent software vendors with porting their applications to AIX.

Ron Barker is a Consulting I/T Specialist in Americas Advanced Technical Support. He has 12 years of experience in AIX and RISC systems. He holds a degree in Communications from Brigham Young University. His areas of expertise include Symmetrical Multiprocessors, Scalable Parallel systems, AIX Version 4 systems administration, and problem determination. He has written or coauthored redbooks, white papers, and presentations on symmetric multiprocessor machines, service processors, and various upgrade and migration topics.

Mauro Minomizaki is an RS/6000 Product Specialist in Brazil. He has three years of experience with AIX and the RS/6000 systems field. He holds a

degree in Data Processing I/T from Sao Paulo Institute of Technology (FATEC-SP). He has worked at IBM for five years. His areas of expertise include RS/6000 systems, RS/6000 SP systems, AIX, performance tuning, sizing and capacity planning, HACMP, HAGEO, and problem determination.

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

Larry Amy, IBM Austin

Barry Barnett, IBM Austin

Bill Beebe, IBM Austin

John Borkenhagen, IBM Rochester

Jayesh Patel, IBM Austin

Salvatore Storino, IBM Rochester

Kurt Szabo, 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 127 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. Overview of the S-Series

The IBM S-Series high-end enterprise server family provides the power, capacity, reliability, and expandability to help move customers into the next generation of mission-critical commercial computing.

The S-Series includes the RS/6000 Model S80 and the RS/6000 Model S70 Advanced. These systems feature the IBM 64-bit symmetric multiprocessing (SMP) server architecture. The Model S80 incorporates IBM state-of-the-art copper chip technology to provide faster more-reliable processors for commercial environments.

The S-Series machines, along with IBM industrial-strength UNIX, AIX Version 4.3, enable customers to manage the evolution of their businesses into 64-bit computing while continuing to run existing applications. The 64-bit hardware addressing capability of the S-Series enables current 32-bit programs to run on systems with very large memory and processor configurations. The S-Series I/O subsystem supports a wide range of 32-bit and 64-bit PCI adapters.

1.1 Series description

Members of the S-Series include the Model S80 and the Model S70 Advanced. They are nearly identical in appearance, consisting of a central electronic complex (CEC) and at least one I/O Drawer, which is housed in a separate I/O Rack. Figure 1 on page 2 shows an example of a Model S80. The CEC is shown on the right, and the I/O Drawer is housed in the I/O Rack on the left.

The CEC and I/O Drawers are connected by Remote I/O (RIO) and System Power Control Network (SPCN) cables. Each I/O Rack can contain 1 or 2 I/O Drawers, up to a maximum of four drawers per system. Each drawer has a total of 14 PCI slots configured across four 33 MHz PCI buses. There are five 64-bit and nine 32-bit slots. In the base configuration, the primary drawer has 11 available slots: One 32-bit slot is required for the service processor, and two 32-bit slots are used for the PCI Ultra SCSI controllers serving the media bays and the first SCSI disk six-pack. To populate the second six-pack, a third PCI Ultra SCSI controller is required.

Both the Model S80 and the Model S70 Advanced can be attached to an RS/6000 SP to provide additional online transaction processing (OLTP) and database capability to the cluster. The S-Series machines can also be configured for high availability using IBM High Availability Cluster Multi

Processing (HACMP) software and redundant hardware components. The S-Series servers attach to the latest IBM storage technology products, giving customers a broad choice of tape and disk storage options. The S-Series supports advanced communications adapters and provides exceptional performance in online transaction processing, Enterprise Resource Planning (ERP), Business Intelligence (BI), Web serving, and e-commerce.

Figure 1. RS/6000 Model S80 Enterprise Server

The S-Series machines are designed for high levels of reliability, availability, and serviceability. Each has a service processor that enables local or remote access to the server while the system is off line. The service processor constantly monitors key system hardware components and, if so configured, can detect operating system hangs and force a restart. The service processor provides early power-off warnings and has facilities for hardware error analysis. It can be configured to dial IBM Service automatically if the system cannot boot. Additionally, IBM makes a program called Service Director for RS/6000 available to warranty and maintenance customers. Service Director monitors a running AIX system, checks for errors, and notifies IBM Service and system administrators of new, impending, or potential problems. Service Director assists customers and IBM Service in achieving higher availability and improved serviceability.

1.1.1 Model S80 specifics

The Model S80 base configuration has a 6-way 450 MHz RS64 III processor card. Each processor has 8 MB of level 2 (L2) cache. Minimum system memory is 2 GB of SDRAM. Cache and system memory is capable of single-bit error correction and double-bit error detection (ECC).

A 6-way system can be expanded to a 12-, 18- or 24-way system, with each additional processor card, called a book, containing six processors. Memory can be increased to 64 GB. Memory must be installed in four-card sets, called quads, which are available in the following sizes: 1 GB (4 x 256 MB), 2 GB (4 x 512 MB), 4 GB (4 x 1024 MB), 8 GB (4 x 2048 MB) and 16 GB (4 x 4096 MB). There are 16 memory book slots, allowing a total of four quads to be installed. Memory cards must be identical within each quad.

The I/O Rack holds the primary I/O Drawer, which contains:

- Service processor
- High-performance disk drive
- 32X Maximum speed CD-ROM
- 1.44 MB 3.5-inch diskette drive
- Two PCI Ultra SCSI controllers

Up to three additional I/O Drawers can be added. Each drawer provides 14 PCI slots across four independent buses. Each drawer may contain two Ultra SCSI hot-pluggable disk 6-packs. Existing RS/6000 7015 Model R00 and 7014 Model S00 racks can also be used for additional rack-mounted storage and communication devices.

A fully-configured system consists of 24 processors, 64 GB of system memory, 53 available PCI adapter slots (the service processor and two Ultra SCSI controllers take up three slots in the first I/O Drawer), 48 hot-pluggable disk bays, and seven available media bays. If the PCI Dual Channel Ultra2 SCSI adapter (#6205) is ordered, it can replace two PCI Single-Ended Ultra SCSI adapters, bringing the maximum number of available PCI slots on a fully-configured system to 54.

PCI adapters may be used to support a wide range of communications and storage subsystems. Supported communications adapters include Gigabit Ethernet, 10/100 Mbps Ethernet, Standard Ethernet, Token Ring, Asynchronous Transfer Mode (ATM), and Fiber Distributed Data Interface (FDDI). Supported storage device protocols include Ultra SCSI, Ultra2 SCSI, Serial Storage Architecture (SSA), and Fibre Channel-Arbitrated Loop (FC-AL).

The Model S80 is shipped and delivered with internal adapters and devices already installed and configured. The Model S80 requires AIX Version 4.3.3 or later, which comes with every system ordered and, if desired, can be pre-installed.

1.1.2 Model S70 Advanced specifics

The base Model S70 Advanced server configuration features one 4-way 262 MHz RS64 II processor book. Each processor has 8 MB of L2 cache. The minimum system configuration includes 1 GB of SDRAM memory.

Two additional processor books may be added to the CEC of the Model S70 Advanced, allowing for 4-, 8- and 12-way configurations. System memory can be expanded to 32 GB. As with the Model S80, memory must be installed in four card sets called quads. The Model S70 Advanced has 20 memory card slots allowing a total of five quads. Memory is available in quads of 1 GB (4 x 256 MB), 2 GB (4 x 512 MB), 4 GB (4 x1024 MB), and 8 GB (4 x 2048 MB).

The I/O Rack holds the primary I/O Drawer, which contains:

- Service processor
- High-performance disk drive
- 32X Maximum speed CD-ROM
- 1.44 MB, 3.5-inch diskette drive
- Two PCI Ultra SCSI controllers

Three additional I/O Drawers may be added. No more than two I/O Drawers may reside in a single rack. External disk drawers may be installed in racks as space permits.

A fully-configured system consists of 12 processors, 32 GB of system memory, 53 available PCI adapter slots, 48 hot-pluggable internal disk bays, and seven available media bays. In the primary I/O Drawer, the service processor uses one slot, and two slots are taken for the SCSI cards that control the media bays and the first disk six-pack. If the PCI Dual Channel Ultra2 SCSI adapter (#6205) is ordered, it can replace two PCI Single-Ended Ultra SCSI adapters, bringing the maximum number of available PCI slots on a fully-configured system to 54.

PCI adapters may be used to support a wide range of communications and storage subsystems. Supported communications adapters include Gigabit Ethernet, 10/100 Mbps Ethernet, Standard Ethernet, Token Ring, ATM and

FDDI. Storage device protocols include Ultra SCSI, Ultra2 SCSI, SSA and FC-AL.

The Model S70 Advanced is shipped and delivered with adapters and devices already installed and configured. System power supplies are designed for maximum configurations. The Model S70 Advanced comes with AIX Version 4.3. The operating system can be preinstalled, if desired.

1.2 Competitive positioning

The performance, expandability, and reliability of the Model S80 and Model S70 Advanced make them strong candidates for the most complex mission-critical e-business applications, as well as ERP, BI, server consolidation, and OLTP solutions. The Model S80 in particular offers leadership performance that makes it a very attractive choice for large commercial and e-business workloads.

The Model S80 is the natural growth path for the Models S70 and S70 Advanced. It offers more performance, memory, and storage. A fully-configured Model S80 is roughly 3.5 times more powerful than a fully-configured Model S70 Advanced.

The S-Series machines are designed from the ground up for mission-critical environments. They come with a service processor, redundant power and cooling, and hot-swap disk drives. You can cluster S-Series models with High Availability Cluster Multi-Processing (HACMP) for AIX to achieve even greater availability and horizontal growth.

The Model S80 has the capacity to meet current and future needs with a choice of 6, 12, 18, or 24 processors and extensive memory options from 2 GB up to 64 GB. The Model S70 Advanced offers excellent price performance with four, eight, or 12 processors and up to 32 GB of memory. Both systems can be configured with up to 53 available PCI expansion slots and up to 45 TB of disk storage.

Competitive systems sold in these application environments include the Sun Enterprise 6500 and 10000 servers, the HP V2500, and possibly some clustered Alpha chip systems from Compaq. The Model S80's leading performance and the excellent price/performance characteristics of the S-Series machines make them very strong contenders in the high-end commercial computing marketplace.

1.3 Performance benchmarks

In July 1999, IBM released some preliminary performance numbers for the then-unannounced Model S80. SPECWeb96 and Java VolanoMark benchmark results were released. Shortly after availability of the Model S80, benchmark results for online transaction processing were released, confirming that the Model S80 is the performance leader among high-end UNIX servers. Changes made to the system for commercial applications, such as doubling the number of processors, almost doubling the clock rate, improving the cache, and increasing processor bandwidth, have produced many new benchmark records.

1.3.1 Online transaction processing

Based on the TPC-C benchmark, the Model S80 is the world's most powerful system for transaction processing. In addition, Transaction Processing Performance Council results show a single Model S80 system with 24 processors can perform more transactions than a cluster of four 24-processor Sun E6500 servers (96 microprocessors in total).

The Model S80 is also the best value among large enterprise servers with the lowest cost per transaction among the 10 most powerful systems. The S80 performed 135,815 transactions per minute (tpmC) at $52.70 per transaction ($/tpmC), which is 50 percent better performance for the dollar than the Sun E10000 and 45 percent better than Sun's cluster result.

For more information on this benchmark result, refer to the Transaction Processing Performance Council Website at the following URL:

`http://www.tpc.org`

1.3.2 SPECWeb96

A new SPECWeb96 record was set with a 12-way Model S80 using a prerelease version of AIX 4.3.3. The Model S80 delivered a SPECweb96 benchmark of 40,161 http operations per second, which is 66 percent better than the 24,139 http operations per second posted by the former leader, an eight-way HP 9000/N4000 server from Hewlett-Packard. SPECWeb96 measures the maximum number of Hypertext Transfer Protocol (HTTP) operations per second. Fast servers that are reliable and can handle large numbers of users are considered critical to companies expanding their businesses and their customer relationships onto the Internet.

For more information on the SPECWeb benchmark, refer to the Standard Performance Evaluation Corporation Website at the following URL:

http://www.spec.org

1.3.3 Java environment performance

The Java VolanoMark record was achieved on a six-way Model S80. It transferred 22,906 messages per second, which is a 92 percent increase over the previous top result posted by Hewlett-Packard with an N Class server configured with eight processors. The VolanoMark is designed to accurately predict real world Java environment performance and scalability. The benchmark creates a large number of active socket connections and threads and forces the system to constantly switch among them all. This behavior is key to successfully deploying Java environments for e-business applications.

While hardware improvements were important in setting these new records, AIX enhancements were also very significant. AIX 4.3.3 offers several Java environment, Web-performance, and scalability enhancements.

1.4 AIX Version 4.3 and commercial computing

AIX Version 4.3.3 offers a number of performance and system administration enhancements for commercial systems.

1.4.1 Kernel scalability enhancements

Kernel scalability enhancements have greatly increased OLTP throughput. Because S-Series machines are capable of handling large memory configurations, AIX 4.3.3 supports multiple lists of free memory frames. A frame is a 4 KB unit of real memory. It maps 1:1 to a 4 KB page of virtual memory. The latest kernel also supports multiple page replacement daemons. These constantly-running processes manage real memory by deciding whether the contents of a location in memory should remain where they are or be moved to disk, where it will take more time to retrieve them in the future. Allowing multiple lists and having multiple daemons reduces memory contention and latency (the time needed to retrieve data or instructions needed by a processor).

In another kernel enhancement, runnable threads are assigned to local run queues on a per-processor basis. This simplifies the dispatcher's decision about which thread to run next by reducing lock contention and eliminating time-consuming calculations needed to maintain affinity between a processor and its cached data. The algorithms used by the dispatcher have been tuned to provide better transaction throughput on busy SMP systems. At AIX 4.3.3, user threads generate less cache interference and maintain a greater affinity to a single processor.

Other AIX 4.3.3 enhancements include:

- The ability to mirror striped logical volumes
- A guarantee that mirrored logical volume copies will exist on different disks
- The ability to do an online backup of a mirrored journaled file system
- The ability to create a system backup image on a Write Once Read Many (WORM) CD drive
- Console messages are logged to a file and time stamped

Faster reboot times may be possible because the system allows some device configuration methods to run in parallel. This option may produce a faster reboot when multiple devices, such as SCSI disks, TTYs, and multiport asynchronous adapters, are connected to a system. A serialization mechanism is used when the configuration manager recognizes that new devices have been added. The system configuration manager still follows the hierarchical device tree, but it can traverse multiple branches at the same time, as long as nothing has been changed from the previous boot. AIX 4.3.3 supports up to 16 concurrent configuration methods.

1.4.2 Workload management

A new facility called workload management comes with AIX 4.3.3. It gives a systems administrator greater control over how the scheduler and Virtual Memory Manager (VMM) determine priorities and allocate system resources.

Although the CPU and memory workload management functions described here can be used with any RS/6000 running AIX 4.3.3, it is expected that they will most commonly be used on large servers running multiple differing workloads that may sometimes compete with one another.

Workload management can help provide isolation between user communities that make very different demands on the system. Some users are likely to run interactive or CPU-intensive applications, while others may require batch or memory-intensive capabilities.

Workload management is fundamentally different from operating system partitioning, which is also called logical partitioning or LPAR. LPAR has been implemented on S/390 and AS/400. Partitioning isolates and dedicates hardware resources to multiple copies of the operating system running on a single machine. In contrast, the goal of workload management is to make more efficient and flexible use of CPUs and memory under the control of a single copy of the operating system.

A more in-depth discussion of workload management can be found in Section 4.2, “Workload management” on page 68.

Chapter 2. Detailed descriptions

This chapter takes an in-depth look at the hardware packaging and features that comprise the RS/6000 Model S80 and Model S70 Advanced Enterprise Servers. The standard and optional features of each model are described along with the features of the I/O Drawer and I/O Rack that are identical on both systems.

This chapter also includes some guidelines for using the PCRS6000 Configurator to create a valid system order.

2.1 Model S80 description

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

The Model S80 is packaged as a central electronic complex and an I/O Rack. The Model S80 entry configuration starts with a six-way scalable SMP system that utilizes the 64-bit, 450 Mhz, RS64 III processor with 8 MB of Level 2 (L2) cache per processor. The 6-way SMP can be expanded to a 24-way SMP, and the system memory can be expanded to 64 GB.

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

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

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

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

2.1.1 Base configuration

A new Model S80 order must include a minimum of the following items:

- One Model S80 system unit, which provides the CEC (Central Electronic Complex) enclosure, system backplane, and CEC power
- Cabling for connection to the primary I/O Drawer
  - Two Remote I/O Cable - CEC to I/O Drawer (#3143 or #3144)
  - One system control and initialization cable (#6000)
  - Two power control cable, Processor Complex to I/O Rack (SPCN)(#6008)
- One Remote I/O Hub - Dual Loop (#6503)
- One 6-way processor complex
  - First RS64 III Processor, 6-Way SMP 450 MHz, 8 MB L2 Cache (#5318)
- 2 GB Memory Minimum - choose from:
  - Two 1024 MB R1 Memory (4 x 256 MB Cards) (#4190)
  - One 2048 MB R1 Memory (4 x 512 MB Cards) (#4191)
  - One 4096 MB R1 Memory (4 x 1024 MB Cards) (#4192)
  - One 8192 MB R1 Memory (4 x 2048 MB Cards) (#4193)
  - One 16384 MB R1 Memory (4 x 4096 MB Cards) (#4194)
- One I/O Rack (#7000)
- One Primary I/O Drawer including:
  - I/O Drawer, 10EIA, AC Power, Unconfigured (#6320)
  - Primary I/O Drawer Group (#6321)
  - Support Processor Group (#6326)
  - 32X (Max) SCSI-2 CD-ROM Drive (#2624)
- 9 GB Hard Disk (Minimum) - choose from:
  - One - 9.1 GB Ultra-SCSI 16-bit Hot-Swap Disk Drive (#2913)
  - One - 9.1 GB 10,000 RPM Ultra-SCSI Hot-Swap Disk Drive (#3008)

- One - 18.1 GB Ultra-SCSI 16-bit Hot-Swap Disk Drive (#3104)

- SCSI Backplanes, Adapters, and Cables. Choose either Single or Dual Backplane Solution.
  - Single Backplane solution, consisting of:
    - One - SCSI 6-Pack Hot-Swap Backplane/Power Cable (#6547)
    - Two - PCI Single-Ended Ultra SCSI Adapter (#6206)
    - One - SCSI-2 Backplane-to-DASD 6-Pack Cable (#2447)
  - Dual Backplane solution, consisting of:
    - Two - SCSI 6-Pack Hot-Swap Backplane/Power Cable (#6547)
    - Three - PCI Single-Ended Ultra SCSI Adapter (#6206)
    - Two - SCSI-2 Backplane-to-DASD 6-Pack Cable (#2447)

  In either solution, two of the PCI Single-Ended Ultra SCSI Adapters may be replaced by a single PCI Dual Channel Ultra2 SCSI Adapter (#6205). However, the Ultra2 adapter will function at Ultra SCSI speed.

2.1.2 Processor

Each Model S80 processor card has six RS64 III processors with the associated L2 cache contained on the card. There are 8 MB of L2 per processor. Each processor card has the six processors on a set of two SMP system buses and that dual bus interface is presented to the S80 backplane. All the processor cards in a system need to use the same type and speed of processor. The Model S80 can accommodate 4 CPU cards. The first CPU card is a Type 1 card (#5318) and is included in a base configuration. Up to three additional Type 2 cards (#5319) may be added.

- (#5318) 450 Mhz Card with 8 MB L2 Cache, First
- (#5319) 450 Mhz Card with 8 MB L2 Cache, Additional

2.1.3 Power

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

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

These additional power components are all installed in the front of the Model S80 CEC. The 1000 Watt AC Power Supply is installed in position P05; the Programmable Power Regulator is installed in position R08, and the Processor Power Regulator is installed in position M07. The locations are shown in Figure 2 on page 15.

2.1.4 Memory

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

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

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

The available memory features are as follows:

- (#4190) 1024 MB Memory (4 x 256 MB Cards)
- (#4191) 2048 MB Memory (4 x 512 MB Cards)
- (#4192) 4096 MB Memory (4 x 1024 MB Cards)
- (#4193) 8192 MB Memory (4 x 2048 MB Cards)
- (#4194) 16384 MB Memory (4 x 4096 MB Cards)

The physical layout of the front of Model S80 CEC is shown in Figure 2 on page 15.

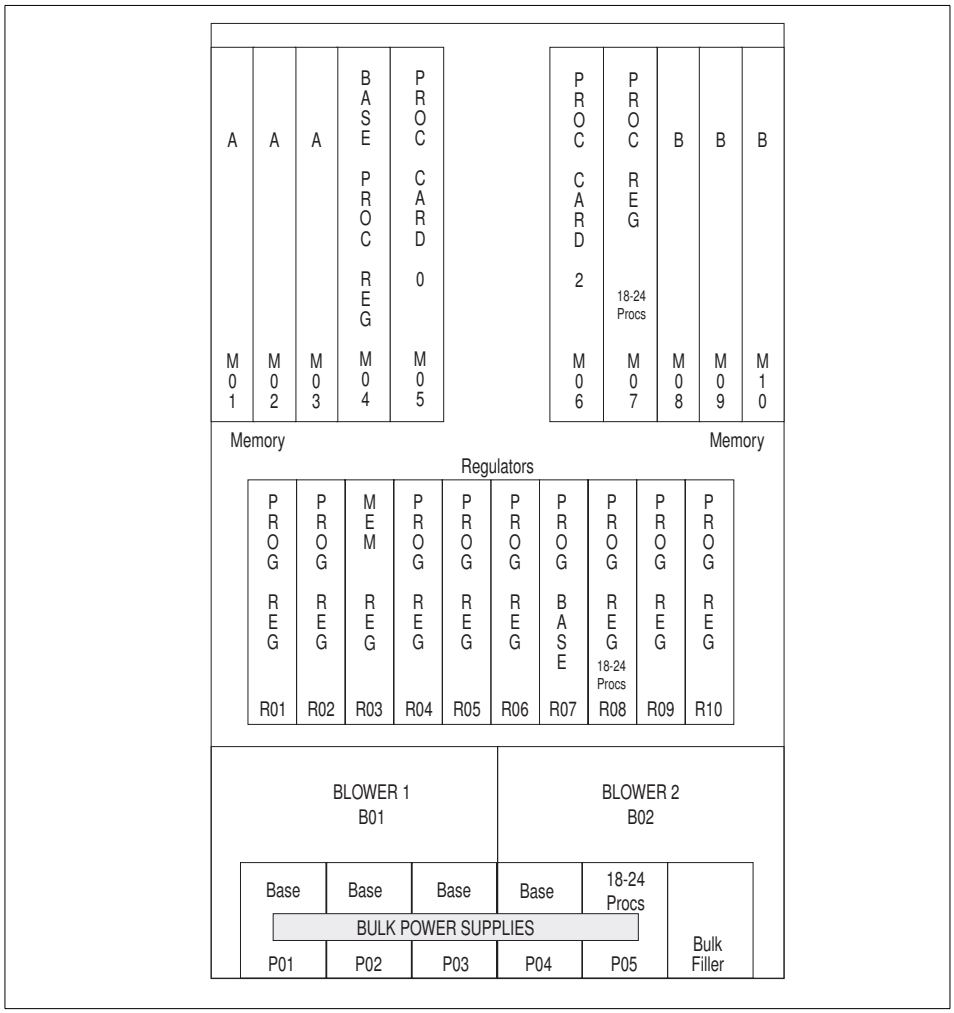

Figure 2. Front view of Model S80 CEC

The memory quads on the system are referred to as A, B, C, and D. Three memory cards, each from quads A and B, are installed on the front of the system. The remaining single card from each quad is installed in the rear of the system. Figure 3 on page 16 shows the rear of the Model S80 CEC. All four cards of each of the C and D quads are installed in the rear of the system.

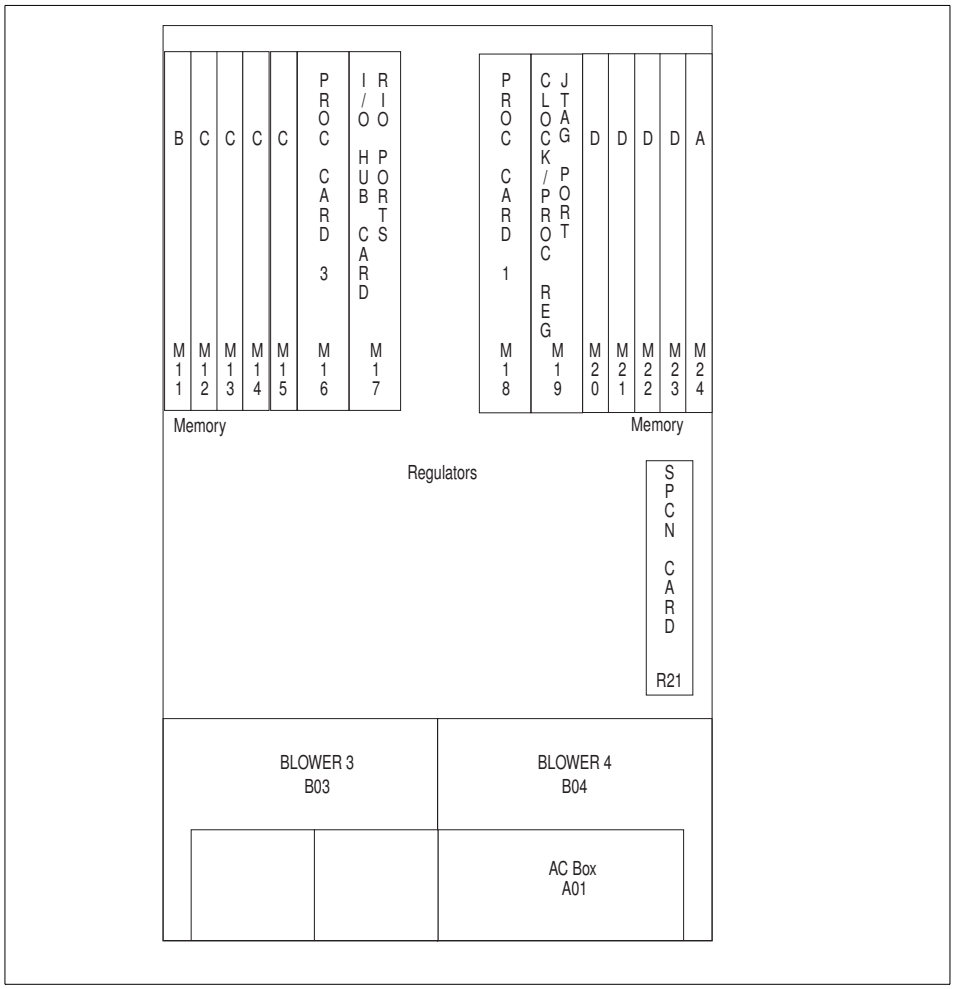

Figure 3. Rear view of Model S80 CEC

2.1.5 Disk drives

Disk drives are available for installation in the I/O Drawers of the Model S80. Disk drives are not supported in the media section of the I/O Drawers. Each I/O Drawer supports up to two SCSI 6-packs. As the name suggests, each 6-pack supports up to six hot-pluggable 16-bit SCSI disk drives. All disk drives must be 16-bit devices. Each 6-pack must be connected to an SCSI adapter using a 6-pack attachment cable.

Disk Drives:

- (#2913) 9.1 GB 1" Ultra SCSI Hot-Swap
- (#3008) 9.1 GB 10,000 RPM Ultra SCSI Hot-Swap
- (#3104) 18.2 GB Ultra SCSI 1" Hot-Swap

SCSI 6-pack:

- (#6547) SCSI 6-pack

6-pack Attachment Cables:

- (#2447) Attachment to PCI SCSI/RAID Adapter

2.1.6 SCSI adapters

Boot support is available from local SCSI Adapters (Ultra SCSI Single-ended Adapter, Differential Ultra SCSI Adapter, Dual Channel Ultra2 SCSI Adapter).

The recommended location for the boot device (SCSI) is within the base I/O Drawer. This configuration provides service personnel with the maximum amount of diagnostic information if the system encounters errors in the boot sequence. The default boot drive is in the lowest location in the six-pack, in the inner-most bay of the I/O Drawer. Manufacturing installs the boot adapter in slot 13. If a boot source other than internal disk is configured, the supporting adapter must also be in the first I/O Drawer.

SCSI Adapters:

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

The base configuration of the Model S80 comes with two Ultra SCSI Single-Ended Adapters (#6206). One adapter drives the media bays, and the other drives the first SCSI 6-pack. These two adapter cards can be replaced by a single Dual Channel Ultra2 SCSI Adapter (#6205), which increases the number of available PCI slots in the primary I/O Drawer. However, the Ultra2 adapter will operate at Ultra speed.

2.1.7 I/O Drawers

See Section 2.3.2, “I/O Drawer description” on page 36, for a full description of the Model S80 I/O Drawer. The original 7 EIA Model S70 I/O Drawers are not supported on the S80.

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

- 1.44 MB Diskette Drive
- (#2624) 32X Speed CD-ROM

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

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

The following lists the feature codes of the drawer groups:

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

2.1.8 Cabling

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

2.1.8.1 System Power Control Network

See Section 2.3.3, “Serial Power Control Network (SPCN)” on page 38, for a full description of the purpose and configuration options for SPCN cables.

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

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

2.1.8.2 Remote I/O cables

See Section 2.3.4, “Remote I/O cabling” on page 40, for a full description of the purpose and configuration options for Remote I/O cables.

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

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

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

2.1.9 Model S80 configurations

Table 1 lists the S80 standard configuration.

Table 1. Model S80 standard configuration

Model S80 Standard Configuration and Features	
Microprocessor	One 6-way 450 MHz RS64 III CPU card
Level 1 (L1) cache	128 KB data/128 KB instruction
Level 2 (L2) cache	8 MB per processor
RAM (minimum)	2 GB
Memory bus width	Quad 512-bit
Ports	One parallel, two serial, one keyboard, and one mouse
Internal disk drive	One 9.1 GB Ultra SCSI (hot-swappable)
Media bays	Two (one available)
Expansion slots	Fourteen PCI (eleven available)
PCI bus width	32- and 64-bit
Memory slots	16
CD-ROM drive	32X (max)
Service processor	Yes
Diskette drive	1.44 MB 3.5-inch diskette drive
SCSI adapters	Two Ultra SCSI PCI adapters
AIX operating system version	Version 4.3 (A one- to two-user server license is included)

Table 2 lists the S80 system expansion capabilities.

Table 2. Model S80 system expansion

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

Model S80 Maximum Configuration	
Internal disk bays	Up to 48 (hot-swappable)
Internal disk storage	Up to 873.6 GB
External disk storage	Up to 38 TB SCSI; up to 45 TB SSA

2.1.9.1 Publications

Table 3 details the publications shipped with the Model S80.

Table 3. Publications shipped with the Model S80

Title	Order Number
RS/6000 Enterprise Server S80 Installation and Service Guide	SA38-0558
RS/6000 Enterprise Server S80 User's Guide	SA38-0557
PCI Adapter Placement Reference	SA38-0538
Diagnostic Information for Multiple Bus Systems	SA38-0509
Adapters, Devices, and Cable information for Multiple Bus Systems	SA38-0516
Customer Installable Options Library CD-ROM[1]	

[1] The CD-ROM is not orderable; it is shipped with the product, and no form number is available.

2.1.10 Model S80 additional features

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

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

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

S Indicates a feature that is supported on the new model during a model conversion; these features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed.

W Indicates features that are supported but will be withdrawn by December 13, 1999 or earlier.

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

Table 4 lists Model S80 optional features and their status.

Table 4. Model S80 optional features

Feature Code	Description	Status
Processors		
5318	RS64 III 6-way 450MHz first	A
5319	RS64 III 6-way 450MHz additional	A
Memory		
4190	1024 MB (4x256)	A
4191	2048 MB (4x512)	A
4192	4096 MB (4x1024)	A
4193	8192 MB (4x2048)	A
4194	16384 MB (4x4096)	A
Host Attachment		
2751	ESCON Control Unit	A
8396	SP Attach Adapter	A
Internal Disk Drives		
2901/9394	4.5 GB Ultra SCSI Hot-Swap	W
2911/3019	9.1 GB Ultra SCSI Hot-Swap	S
2913	9.1 GB 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
3008	9.1 GB 10K RPM Ultra SCSI Hot-Swap	A
Internal Tape Drives		
6142	4/8 GB 4 mm	W
6147	5/10 GB 8 mm	W
6154	20/40GB 8 mm (White)	S
6156	20/40 GB 8 mm (Black)	A
6159	12/24 GB 4 mm	A
Internal CD-ROMs		

Feature Code	Description	Status
2619	20X Speed CD-ROM	S
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2838	GXT120P	W
2830	GXT130P	A
SCSI Adapters		
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
6208	SCSI 2 Fast / Wide	S
6205	Dual Channel Ultra2 SCSI	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	W
6222	SSA Fast-Write Cache Option	W
6225	IBM Advanced Serial RAID	A
6235	IBM Advanced Serial RAID Cache Option	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
2948	ARTIC960Hx 4-Port T1/E1 PCI	W
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
Cryptographic Adapters		
4758-001	Cryptographic Coprocessor (PCI)	S
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A

Feature Code	Description	Status
2988	Turboways 155 PCI MMF ATM	A
Token-Ring Adapters		
2920	Token-Ring Adapter	W
2979	Auto LANStreamer Token-Ring	S
4959	Token-Ring PCI Adapter	A
Ethernet Adapters		
2968	IBM 10/100 Mbps Ethernet	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2986	Fast Etherlink XL 3Com	S
2987	Ethernet AUI / RJ-45	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
Fiber Channel Adapters		
6227	Gigabit Fibre Channel Adapter[1]	A
FDDI Adapters		
2741	SysKonnect SK-NET FDDI-LP SAS	A
2742	SysKonnect SK-NET FDDI-LP DAS	A
2743	SysKonnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A

[1] See Section 2.5.2, “Configuring S-Series servers” on page 48.

2.2 Model S70 Advanced description

The S70 Advanced was announced on October 23rd, 1998 and is an enhanced version of the Model S70. The S70 Advanced can be used as a stand-alone server, but it can also be attached to the RS/6000 SP as an SP-attached server. Using a specially-designed adapter (separately

available) that fits within the S70 Advanced I/O Drawer, the system is capable of connecting directly into the SP Switch fabric. The system is then ideally suited to handle large database transactions while allowing the other SP nodes to act as application servers. When configured as an SP-attached server, the S70 Advanced is managed and controlled just like a regular SP node using the SP-unique Parallel Systems Support Programs (PSSP) systems management software.

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

2.2.1 Processor

The RS/6000 Model S70 Advanced server base configuration features one 4-way 262 MHz RS64 II book, and it can accommodate up to two additional 4-way processor books. The processors include an L1 Cache split into a 64 KB instruction cache and a 64 KB data cache. There is 8 MB of ECC L2 cache for each 262 MHz processor. The following lists the available processor cards:

- (#5313) 262 MHz Card with 8 MB L2 Cache, Left
- (#5312) 262 MHz Card with 8 MB L2 Cache, Right

2.2.2 Power

The Model S70 Advanced base configuration is equipped with sufficient power supplies and regulators. You do not need to order additional components when ordering a fully-configured system.

2.2.3 Memory

The base configuration includes 1 GB of SDRAM-based memory consisting of two 512 MB features. The maximum configuration is 32 GB. Memory sizes cannot be mixed within each four-card set. There is a maximum of 20 memory card slots, which means a maximum of five memory quads.

System memory is accessed through two related but distinct ports. Balanced accesses use both ports in a coordinated parallel manner and can obtain up to twice the data in the same amount of time. Unbalanced configurations function properly, but the memory is only accessed through one port and does not make use of the full memory bus bandwidth. For example, a system with 2 GB of memory will perform better with two 1 GB features installed than if one 2 GB feature is installed.

The following is a list of the available memory features:

- (#4171) 512 MB memory (4x128 MB cards)
- (#4173) 1024 MB memory (4x256 MB cards)
- (#4175) 2048 MB memory (4x512 MB cards)
- (#4177) 4096 MB memory (4x1024 MB cards)
- (#4179) 8192 MB memory (4x2048 MB cards)

The quads on a Model S70 Advanced are referred to as A, B, C, D, and E. The A, B, and C quads are installed on the front of the system CEC. Each quad has two cards on the left side of the system and two cards on the right. A diagram of the CEC front layout is shown in Figure 4 on page 26.

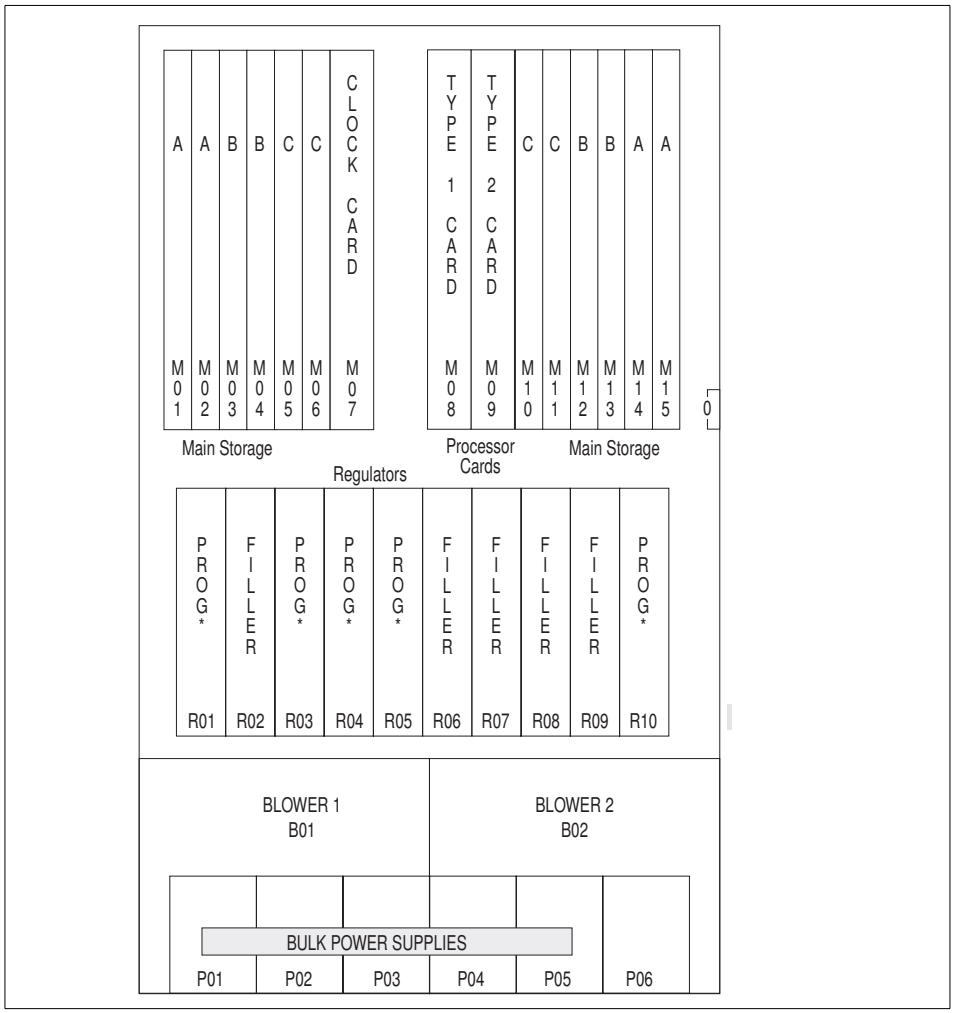

Figure 4. Front view of Model S70 Advanced CEC

The memory cards of the D and E quads are installed in the rear of the system CEC. This is shown in Figure 5 on page 27.

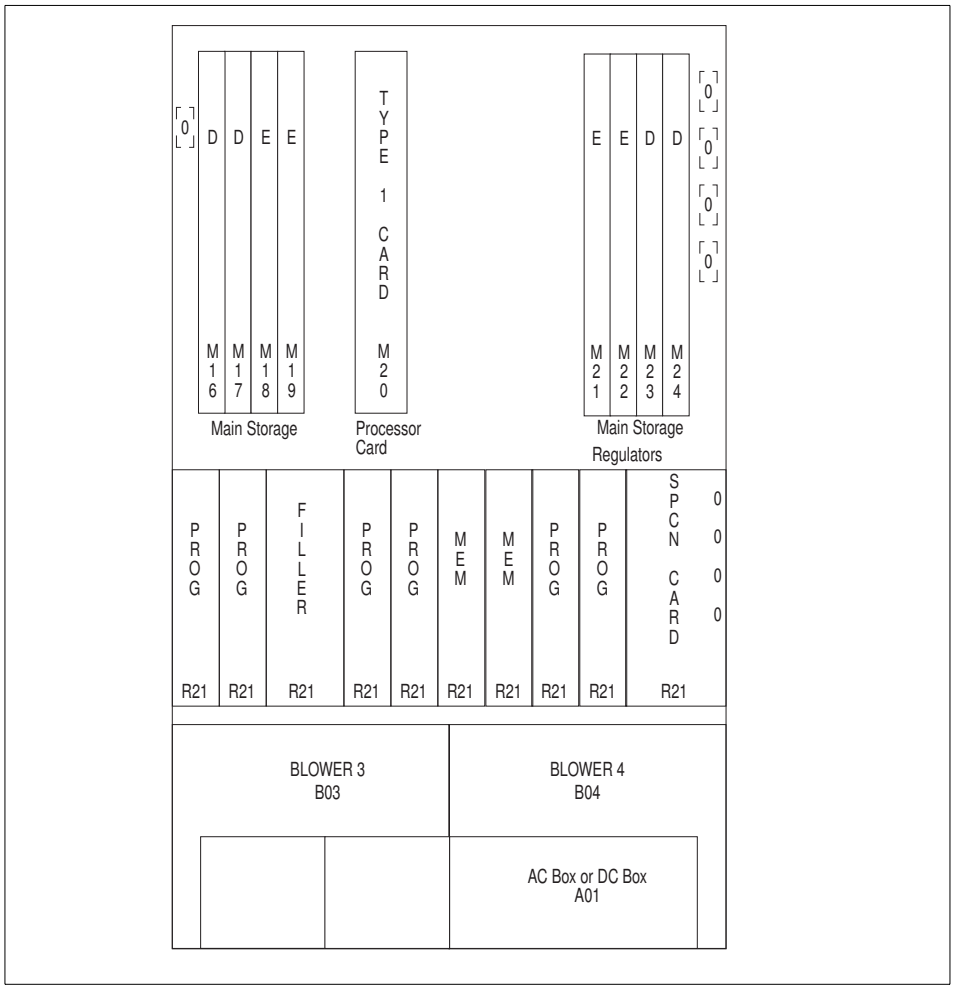

Figure 5. Rear view of Model S70 Advanced CEC

As with the A, B, and C quads, two cards from each quad are mounted on the left side of the CEC, and two cards are mounted on the right side.

2.2.4 Disk drives

Disk drives are not supported in the media section of the I/O Drawers. All disk drives must be 16-bit devices. Note that the (#2901) 4.5 GB Ultra SCSI Fast/Wide feature will be withdrawn as of December 13th 1999.

- (#3104) 18.2 GB 1" Ultra SCSI Hot-Swap
- (#3008) 9.1 GB 10K RPM Ultra SCSI Hot-Swap

- (#2913) 9.1 GB 1" Ultra SCSI Hot-Swap
- (#2901) 4.5 GB Ultra SCSI Fast/Wide

SCSI 6-pack:

- (#6547) SCSI 6-pack

6-pack Attachment Cables:

- (#2447) Attachment to PCI SCSI Adapter

2.2.5 SCSI adapters

Boot support is available from local SCSI Adapters (Ultra SCSI Single-ended Adapter, PCI SCSI-2 Fast/Wide Differential Adapter, PCI Differential Ultra SCSI Adapter, and Dual Channel Ultra2 SCSI Adapter).

The recommended location for the boot device (SCSI) is within the base I/O Drawer. This configuration provides service personnel the maximum amount of diagnostic information if the system encounters errors in the boot sequence. The default boot drive is in the lowest location in the six-pack, in the inner-most bay of the I/O Drawer. Manufacturing installs the boot adapter in slot 13. If a boot source other than internal disk is configured, the supporting adapter must also be in the first I/O Drawer.

- (#6205) Dual Channel Ultra2 SCSI Adapter
- (#6206) Ultra SCSI Adapter
- (#6207) Ultra SCSI Differential Adapter

2.2.6 Cabling

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

2.2.6.1 System Power Control Network (SPCN)

See Section 2.3.3, “Serial Power Control Network (SPCN)” on page 38, for a full description of the purpose and configuration options for SPCN cables. The base configuration includes two SPCN cables for connecting the primary I/O Drawer. Additional I/O Drawers must be cabled using the 15-meter rack-to-rack cable or the 2-meter cable, which can only be used to interconnect two I/O Drawers in the same rack.

The available SPCN cables for Model S70 Advanced systems are as follows:

- (#6006) 2-meter Drawer-to-Drawer Control Cable

- (#6007) 15-meter Rack-to-Rack Control Cable

2.2.6.2 Remote I/O cables

See Section 2.3.4, “Remote I/O cabling” on page 40, for a full description of the purpose and configuration options for Remote I/O cables.

RIO cables are available in two different lengths. The base system order includes two 6-meter cables for connecting the primary I/O Drawer. Additional I/O Drawers must be cabled using the 15-meter rack-to-rack cable or the 2-meter cable, which can only be used to interconnect two I/O Drawers in the same rack. Manufacturing will determine the placement and cabling of I/O Drawers based on the quantity of I/O Racks and RIO cables ordered.

The following remote I/O cables are available as additional features on Model S70 Advanced systems:

- (#3126) 2-meter Remote I/O Cable, Drawer-to-Drawer
- (#3127) 15-meter Remote I/O Cable, Rack-to-Rack

2.2.7 Model S70 Advanced configurations

Table 5 lists the S70 Advanced standard configuration.

Table 5. Model S70 Advanced standard configuration

Model S70 Advanced Standard Configuration and Features	
Microprocessor	One 4-way 262 MHz PowerPC RS64 II CPU card
Level 1 (L1) cache	64 KB data/64 KB instruction
Level 2 (L2) cache	8 MB per processor
RAM (minimum)	1 GB
Memory bus width	Dual 512-bit
Ports	One parallel, two serial, one keyboard, and one mouse
Internal disk drive	One 9.1 GB Ultra SCSI (hot-swappable)
Media bays	Two (one available)
Expansion slots	Fourteen PCI (eleven available)
PCI bus width	32- and 64-bit
Memory slots	Twenty

Model S70 Advanced Standard Configuration and Features	
CD-ROM Drive	32X (max)
Service Processor	Yes
Diskette Drive	1.44 MB 3.5-inch diskette drive
SCSI Adapters	Two Ultra SCSI PCI Adapters
AIX operating system version	Version 4.3 (one to two user server license is included)

Table 6 lists the S70 Advanced system expansion capabilities.

Table 6. Model S70 Advanced system expansion

Model S70 Advanced Maximum Configuration	
SMP configurations	Up to two additional 4-way processor books
RAM	Up to 32 GB
Internal PCI slots	Up to 56 per system (53 available)
Internal media bays	Up to eight per system
Internal disk bays	Up to 48 (hot-swappable)
Internal disk storage	Up to 436.8 GB
External disk storage	Up to 38 TB SCSI; up to 45 TB SSA

Note

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

2.2.7.1 Publications

Table 7 provides details of the publications shipped with the Model S70 Advanced.

Table 7. Publications shipped with the Model S70 Advanced

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

Title	Order Number
PCI Adapter Placement Reference	SA38-0538
Diagnostic Information for Multiple Bus Systems	SA38-0509
Adapters, Devices, and Cable information for Multiple Bus Systems	SA38-0516
Customer Installable Options Library CD-ROM[1]	

[1] The CD-ROM is not orderable; it is shipped with the product. No form number is available.

2.2.8 Model S70 Advanced additional features

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

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

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

S Indicates a feature that is supported on the new model during a model conversion; these features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed.

W Indicates features that are supported but will be withdrawn by December 13, 1999 or earlier.

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

Table 8 lists S70 Advanced optional features and their status.

Table 8. Model S70 Advanced optional features

Feature Code	Description	Status
Processors		
5317	RS64 II 4-way 262 MHz exchange left	A
5316	RS64 II 4-way 262 MHz exchange right	A
5313	RS54 II 4-way 262 MHz optional left	A
5312	RS64 II 4-way 262 MHz optional right	A

Feature Code	Description	Status
Memory		
4171	512 MB (4x128)	A
9168	Base 512 MB (4x128)	S
4173	1024 MB (4x256)	A
4174	1024 MB (4x256) select	S
4175	2048 MB (4x512)	A
4176	2048 MB (4x512) select	S
4177	4096 MB (4x1024)	A
4178	4096 MB (4x1024) select	S
4179	8192 MB (4x2048)	A
4180	8192 MB (4x2048) select	S
Host Attachment		
2751	ESCON Control Unit	A
8396	SP Attach Adapter	A
Internal Disk Drives		
2901/9394	4.5 GB Ultra SCSI Hot-Swap	W
2911/3019	9.1 GB Ultra SCSI Hot-Swap	S
2913	9.1 GB 1" Ultra SCSI Hot-Swap	A
3104	18.2 GB 1" Ultra SCSI Hot-Swap	A
3008	9.1 GB 10K RPM Ultra SCSI Hot-Swap	A
Internal Tape Drives		
6142	4/8 GB 4 mm	W
6147	5/10 GB 8 mm	S
6154	20/40GB 8 mm (White)	S
6156	20/40 GB 8 mm (Black)	A
6159	12/24 GB 4 mm	A
6160	9-Track 1/2" Tape Drawer	W

Feature Code	Description	Status
Internal CD-ROMs		
2619	20X Speed CD-ROM	S
2624	32X Speed CD-ROM	A
Graphics Accelerators		
2838	GXT120P	W
2830	GXT130P	A
SCSI Adapters		
6206	Ultra SCSI SE	A
6207	Ultra SCSI Differential	A
6208	SCSI 2 Fast / Wide	S
6209	SCSI 2 Fast / Wide Differential	S
2493	SCSI 2 Fast / Wide RAID	S
6205	Dual Channel Ultra2 SCSI	A
SSA Adapters		
6215	SSA Multi-Initiator / RAID EL	W
6222	SSA Fast-Write Cache Option	W
6225	IBM Advanced Serial RAID	A
6235	IBM Advanced Serial RAID Cache Option	A
Async Adapters		
2943	8-Port Async EIA-232/422	A
2944	128-Port Async Controller	A
ARTIC Adapters		
2947	ARTIC960Hx 4-Port Selectable	A
2948	ARTIC960Hx 4-Port T1/E1 PCI	W
Digital Trunk Adapters		
6310	ARTIC960RxD Quad Digital	A
Cryptographic Adapters		

Feature Code	Description	Status
4758-001	Cryptographic Coprocessor (PCI)	S
ATM Adapters		
2963	Turboways 155 PCI UTP ATM	A
2988	Turboways 155 PCI MMF ATM	A
Token-Ring Adapters		
2920	Token-Ring Adapter	W
2979	Auto LANStreamer Token-Ring	S
4959	Token-Ring PCI Adapter	A
Ethernet Adapters		
2968	10/100 Mbps	A
2969	Gigabit SX	A
2985	Ethernet BNC / RJ-45	A
2986	Fast Etherlink XL 3Com	S
2987	Ethernet AUI / RJ-45	A
WAN Adapters		
2962	2-Port Multiprotocol PCI	A
Fiber Channel Adapters		
6227	Gigabit Fibre Channel Adapter[1]	A
FDDI Adapters		
2741	SysKonnect SK-NET FDDI-LP SAS	A
2742	SysKonnect SK-NET FDDI-LP DAS	A
2743	SysKonnect SK-NET FDDI-UP SAS	A
ISDN Adapter		
2708	Eicon ISDN DIVA PRO 2.0 PCI S/T	A

[1] See Section 2.5.2, “Configuring S-Series servers” on page 48.

2.3 Remote I/O Subsystem description

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

2.3.1 I/O Rack description

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

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

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

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

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

- A maximum of four I/O Racks can be ordered as features (#7000) per system order. I/O Drawers should be spread between I/O Racks but not exceed two per rack.
- I/O Racks ordered as feature numbers of the Model S70 Advanced or Model S80 (#7000) must contain an I/O Drawer. If additional external communication and storage devices, such as 7133 and 7027, do not fit in the space remaining in the I/O Racks, additional empty I/O Racks should be ordered. The additional I/O Racks should be ordered as products (7014-S00 or 7015-R00) rather than features of the S-Series server. There is no limit to the quantity of 7014-S00 and 7015-R00.
- If the quantity of power distribution units (#6171 or #6173) is greater than one, the space available in the rack is reduced by one EIA.
- The 3570-B11 and 3570-B12 can only be installed in a specific location in the I/O Rack due to drawer interference. Ordering a machine with rack

contents specify feature code #0166 or #0167 will have the base SCSI I/O Drawer installed in a lower location in the rack leaving a nine EIA space empty at the top of the rack. The 3570 is field installed.

- The 3570-C11 and C12 can be installed in any location within the I/O Rack.
- Many 3490 and 3590 tape libraries cannot be installed in S-Series servers' I/O Racks due to interference with the rack door. These machine types are supported for attachment to the S-Series servers, but they must be in a separate rack (7015-R00) unless the I/O Rack door is removed. Instructions for proper removal of the rack door are included in the tape library installation instructions.
- Two IBM 3590-B11 or 3590-E11 tape drives may share a twelve EIA space in the rack. (Order one placement for each pair of units.)
- The 3490 Model F11 and 3590 Model B1A can be installed in the I/O Rack without interference with the rack door installed.

2.3.2 I/O Drawer description

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

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

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

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

Figure 6 shows a 10 EIA I/O Drawer when viewed from the front.

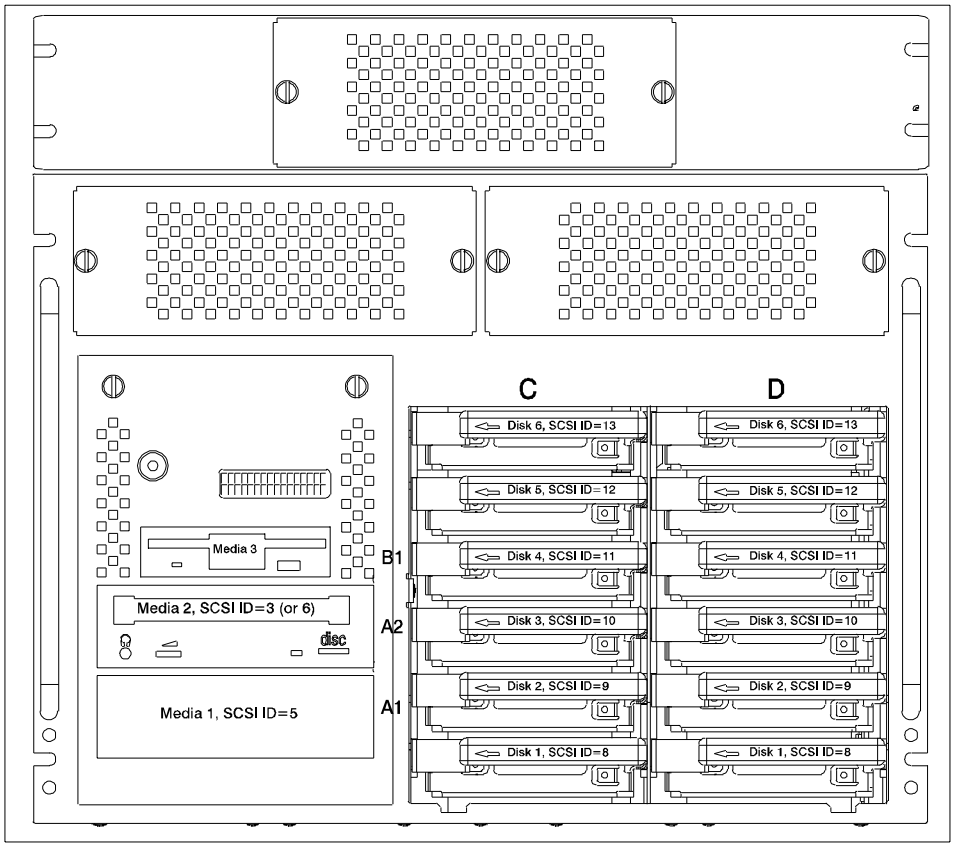

Figure 6. Front view of 10 EIA Drawer

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

The I/O Drawer has redundant power and cooling. If the I/O Racks are ordered with an additional PDU, each drawer can be connected to two PDUs,

thus, making the I/O system highly available from a power supply point of view.

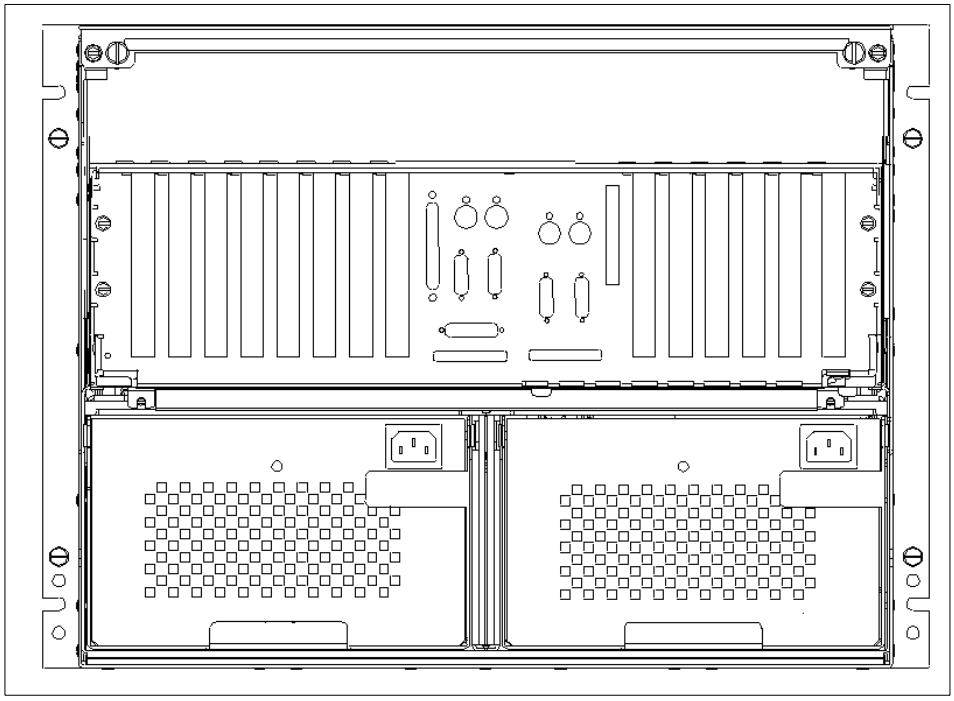

Figure 7. Rear view of 10 EIA I/O Drawer

2.3.3 Serial Power Control Network (SPCN)

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

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

The rules for cabling the SPCN are the same for all S-Series servers. A minimum of two SPCN cables is required for attachment of the first I/O Drawer on each system. Each additional I/O Drawer requires one additional SPCN cable for loop attachment. The available SPCN cable lengths depend on the Model of S-Series server. Figure 8 on page 39 shows the minimal SPCN loop configuration.

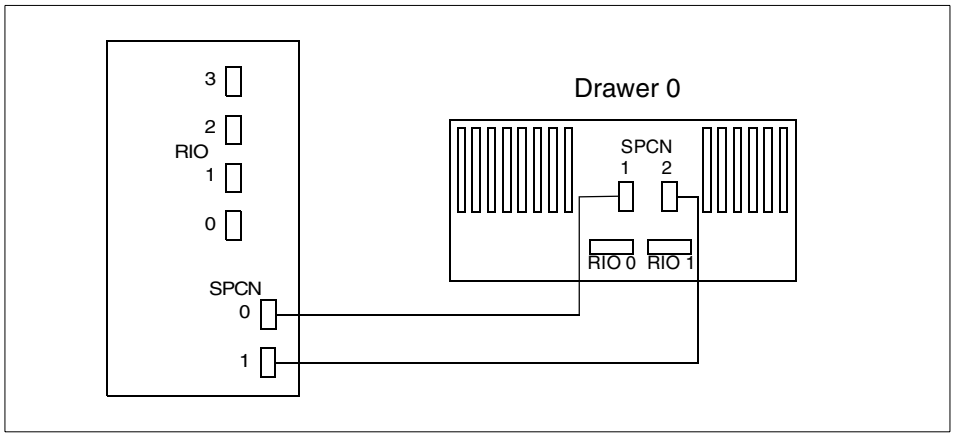

Figure 8. Base SPCN loop with one I/O Drawer

Subsequent I/O Drawers are added to the SPCN loop. The length of SPCN cable chosen will depend on whether the drawer being added is in the same I/O Rack as an existing drawer. Figure 9 shows the maximum SPCN configuration.

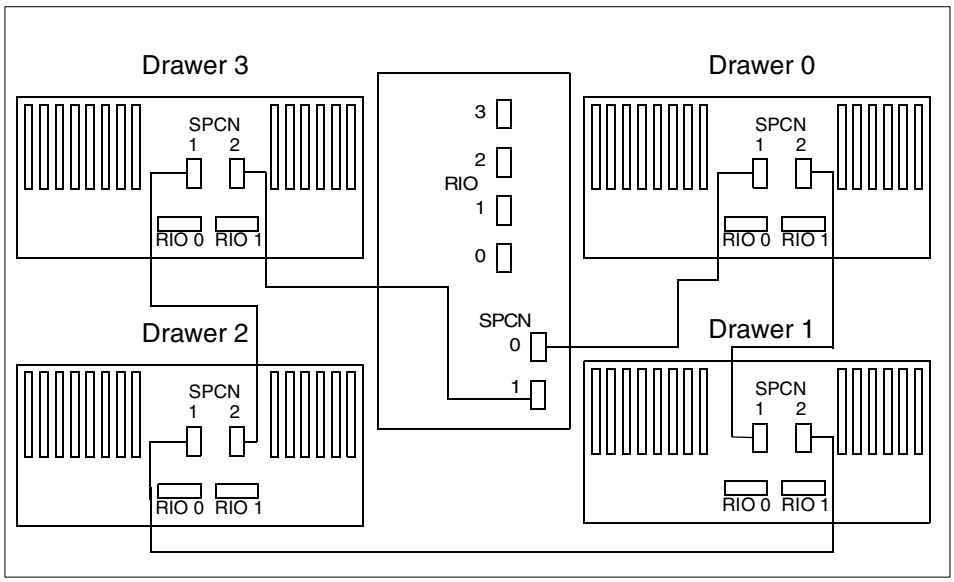

Figure 9. Fully configured SPCN loop with four I/O Drawers

2.3.4 Remote I/O cabling

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

RIO connections operate at 250 MB/s in each direction, 500 MB/s duplex. The loop connection provides redundant paths; so, if a failure occurs in part of a cable, a warning message will be displayed, but the system will continue to operate.

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

The primary I/O Drawer must be installed and connected to RIO port 0 of the CEC. The connection must be made from RI0 port 0 of the CEC to RIO port 0 of the primary I/O Drawer. This connection is required to make the primary drawer the first drawer in the loop that allows the firmware to initialize the system. It also allows the system to find the boot device.

An RIO loop supports up to two drawers. Each S-Series server supports up to two RIO loops. The configuration for a single I/O Drawer is very simple. There can only be one RIO loop since there is only one I/O Drawer. An example of this is shown in Figure 10 on page 41.

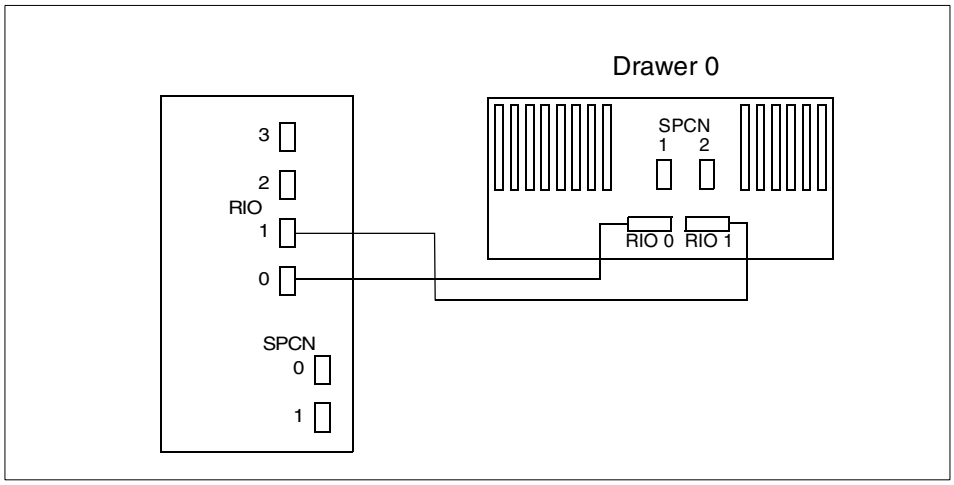

Figure 10. RIO loop with one I/O Drawer

When adding a second drawer, you can cable the additional drawer into the existing RIO loop. This requires one additional RIO cable. An example of this is shown in Figure 11.

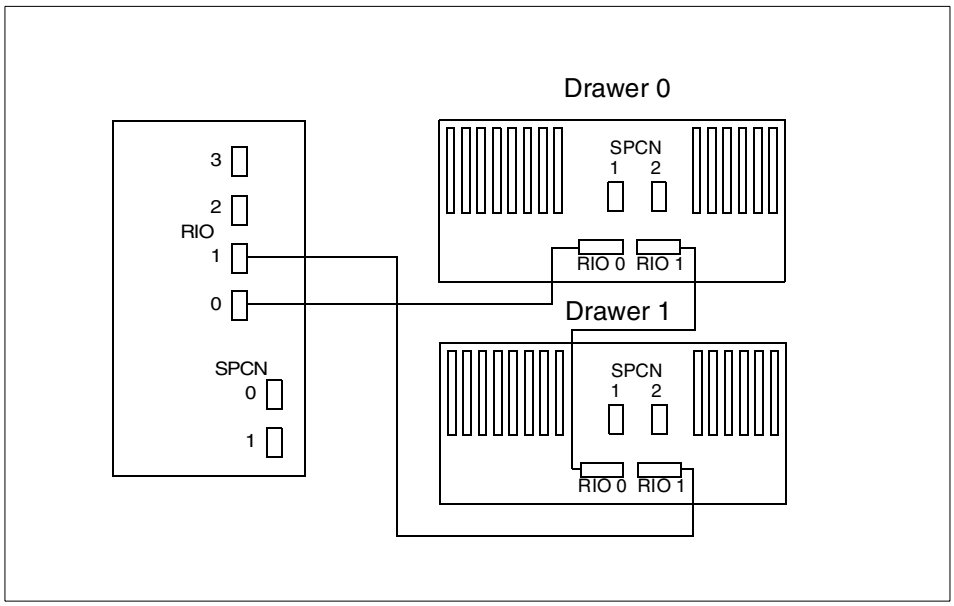

Figure 11. One RIO loop attached to two I/O Drawers

Alternatively, you can cable the second I/O Drawer on its own loop. This option requires two additional RIO cables and provides each drawer with 500 MB of bandwidth. An example of this is shown in Figure 12.

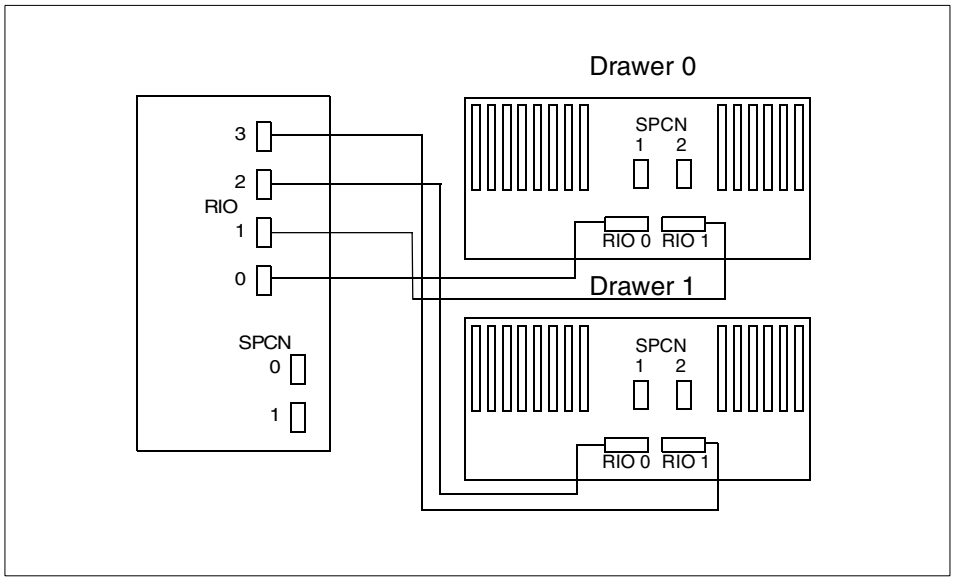

Figure 12. Two RIO loops attached to two I/O Drawers

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

A fully-configured four drawer system must have two RIO loops. An example of this is shown in Figure 13 on page 43.

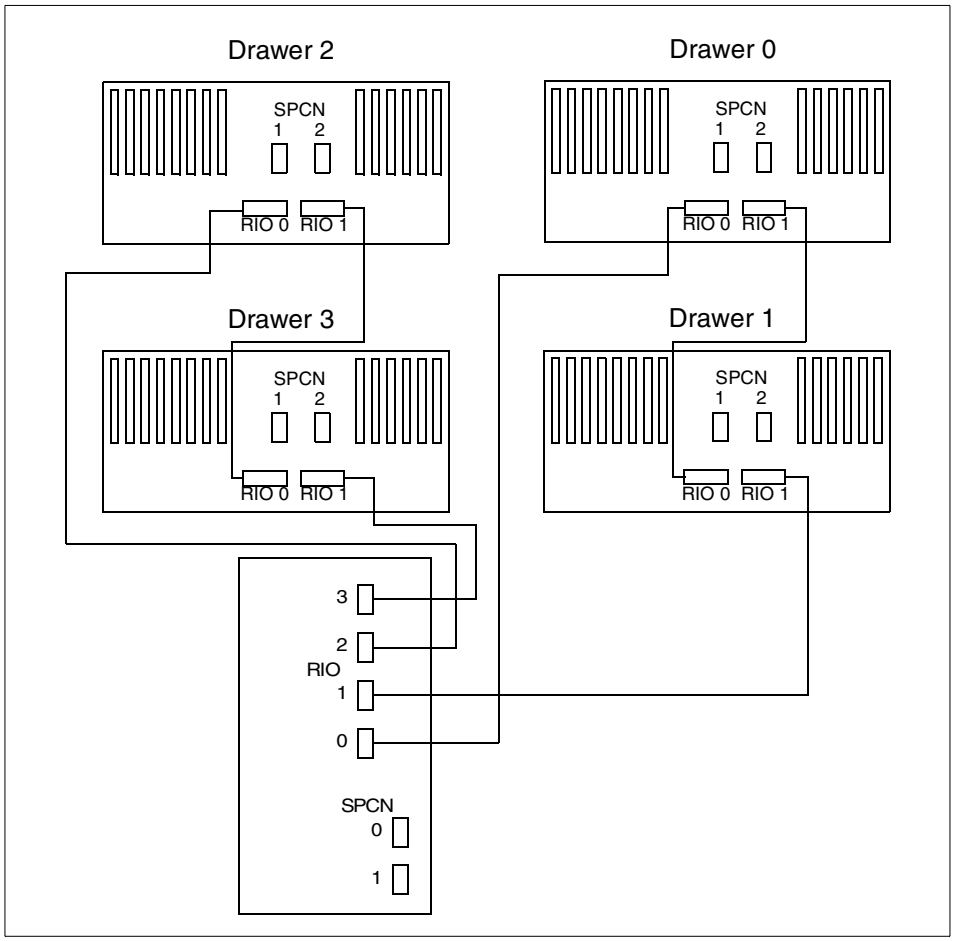

Figure 13. Two RIO loops attached to four I/O Drawers

2.4 Upgrade paths to Model S80

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

The model upgrade consists of the following items:

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

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

- One - First RS64 III Processor, 6-Way SMP 450 MHz, 8 MB L2 Cache (#5318)
- 2 GB Minimum System Memory (See Section 2.4.2, “S70 Advanced to S80 memory conversion” on page 45)
- Cabling for connection to the primary I/O Drawer includes:
  - One - System Control And Initialization Cable - Upgrade Indicator (#8006)
  - Two - Power Control Cable - Upgrade Indicator (#8008)
  - Two - Remote I/O Cables (#3143 or #3144)
- One - Remote I/O Hub - Dual Loop - (#6503)
- One - Support Processor Group - Model Upgrade Only (#8326)

Note

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

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

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

2.4.1 S70 Advanced to S80 processor conversion

Model S70 Advanced systems that are converted to Model S80 systems will require replacement of all system processors. The first Model S70 Advanced

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

Table 9 shows the processor feature conversions that are available to Model S70 Advanced customers who convert their systems to Model S80 systems:

Table 9. S70 Advanced to S80 processor conversion

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

2.4.2 S70 Advanced to S80 memory conversion

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

The following memory feature conversions are available to Model S70 Advanced customers who convert their systems to Model S80 systems:

Table 10. S70 Advanced to S80 memory conversion

From	To
1024 GB	
(#4173) 1024 MB R1 Memory	(#4190) 1024 MB Memory
(#4174) 1024 MB R1 Memory Select	(#4190) 1024 MB Memory
2048 MB	

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

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

2.5 Using the configurator

This section contains a brief explanation of the PCRS6000 Configurator as well as some advice about configuring S-Series machines, both as stand-alone servers and SP-node attachments.

2.5.1 What is the PCRS6000 Configurator?

The PCRS6000 Configurator is an application that provides configuration support for hardware, software, and peripherals associated with the IBM RS/6000 product line and available for marketing.

Functions provided by PCRS6000 include:

- The ability to create, save, restore, and print Hardware/Software, Hardware Only, Peripheral Only, or Software Only configurations.
- Hardware/Software interaction for identifying prerequisite and incompatibility conditions.
- Iterative support for reentering product categories and continuous modification and adjustments to the configuration.
- The ability to modify new or existing Initial Order, MES, or upgrade configurations.
- The ability to modify an installed base prior to beginning MES or upgrade configuration.

- Support for Feature Exchange and Feature Conversion.
- The ability to download and upload saved files to the Host/IBMLink.
- Configuration, price, and price total information is displayed as the configuration is built.
- The ability to apply discounts by percentage to individual products or the total configuration.
- Limited native language support is provided for product descriptions and features (French, German, Italian, Portuguese, Spanish, and Swiss-English).

Figure 14 shows a screenshot of the PCRS6000 configurator.

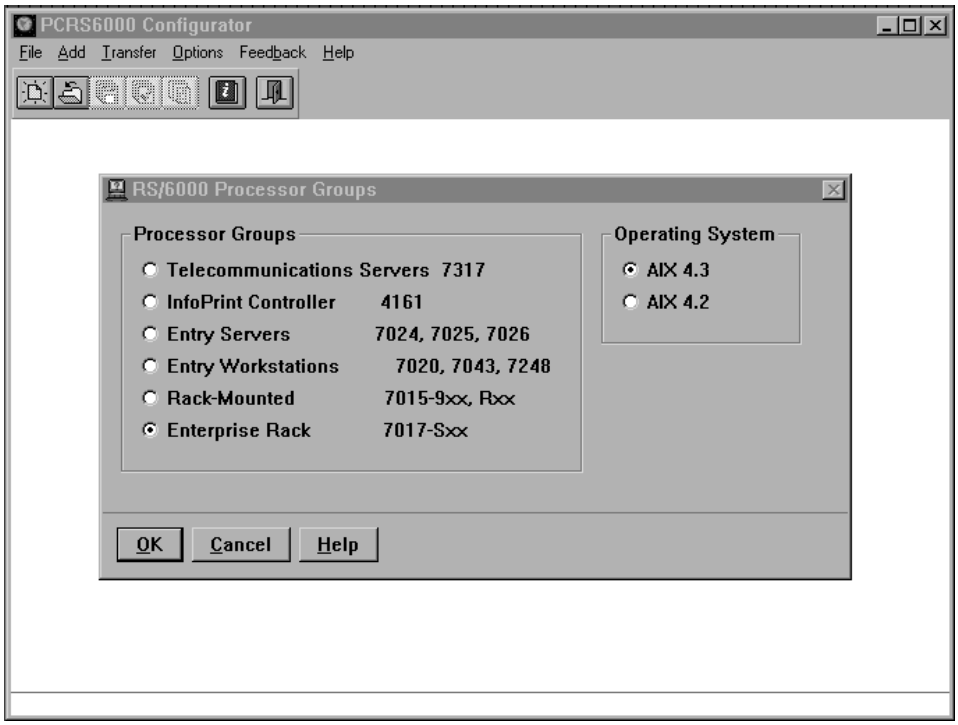

Figure 14. PCRS6000 Configurator

2.5.2 Configuring S-Series servers

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

- In an HACMP configuration, the standard native serial ports on the S-Series machines are not available for heartbeat cabling; therefore, you

must install either the 8-port asynchronous adapter (#2943) with serial-to-serial cable (#3125) or the 128-port Asynchronous controller (#2944) with 128-port Asynchronous cable (#8136), rack-mountable Remote Asynchronous Node (#8136), RJ-45 to DB-25 converter cable (#8133), and serial-to-serial cable (#3125).

- Either an IBM-supported ASCII terminal with an attachment cable or an IBM-supported graphics display with an attachment cable and graphics accelerator is required for initial setup and must be available locally for service. OEM asynchronous terminals are not recommended since they may not transmit the same character sequence and, thus, will not be recognized by the Service Processor.
- SP Node Attachment configurations are described in Section 2.5.3, “Special considerations for SP external node attach” on page 51.
- Gigabit Fiber Channel Adapter (#6227) is available on an MES/Upgrade basis only.

2.5.2.1 Two S-Series Servers sharing an I/O Rack

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

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

When placing an order for multiple systems, the first system should be configured as normal. 10 EIA units of space in the I/O Rack should be reserved for the I/O Drawer of the second system. This is achieved by selecting feature code #0177 as a rack contents specify feature.

When configuring the second system, there will be a feature change that indicates a rackless system order. The feature is selected in the rack options configuration screen, as shown in Figure 15.

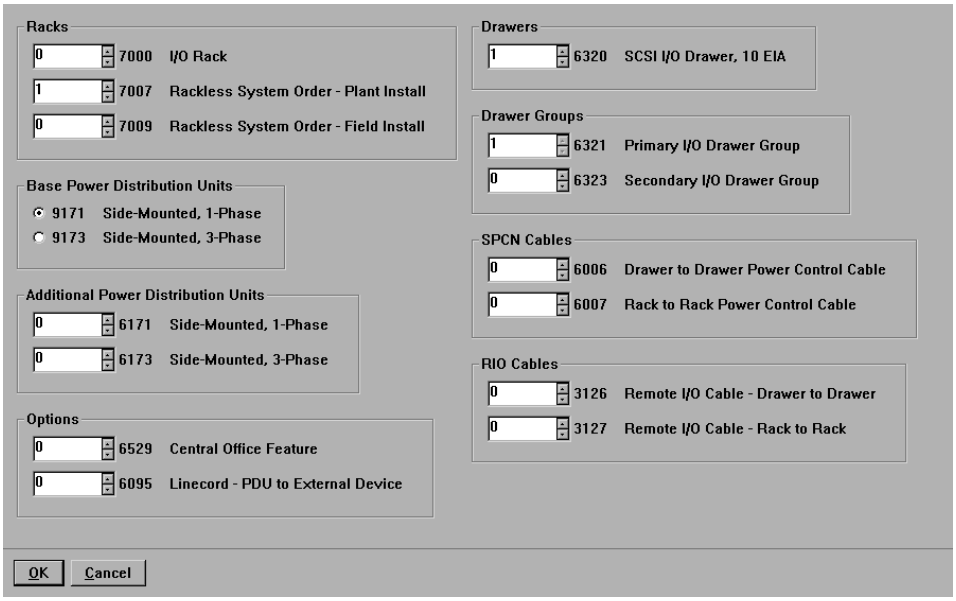

Figure 15. Rack options dialog box

There are two ways of configuring a rackless system:

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

Note

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

2.5.3 Special considerations for SP external node attach

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

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

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

Each S-Series server system that is to function as an attached SMP server within the IBM RS/6000 SP environment must have a minimum of one ethernet adapter. This adapter must be recognized by the system as *en0* and must reside in slot #5 of the primary I/O Drawer.

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

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

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

If this is to be a switchless attachment to the RS/6000 S80 and RS/6000 S70 Advanced servers, then the RS/6000 SP system must also be switchless.

There must be available (unused) switch node numbers in your SP System Data Repository even though the switch is not being used. That means an RS/6000 S80 or RS/6000 S70 Advanced server cannot be attached to a single frame system where all 16 nodes are occupied.

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

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

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

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

- (#9122) NODE ATTACHMENT FEATURE - An order for this feature will generate an order for the two 15 Meter RS-232 cables for hardware control and S1TERM connectivity between the CWS and the SP-attached server. It also will generate an order for a 10 meter ground cable. It also traps some data so that the RS/6000 SP system configurator session can keep track of how many nodes (real and logical) are in the system.
- (#9310) SWITCH CONNECTION CABLE - This feature is required only if the SP-attached server is switch attached. It results in the ordering of one 10 meter switch connection cable. The 10 meter cable is the only supported length at this time.
- (#9123) FRAME ATTACHMENT FEATURE - This feature keeps track of how many frames are currently in your RS/6000 SP system. Since the S-Series SP-attached server is both a logical node and a logical frame in the PSSP code logic, it is important to track this information to avoid exceeding allowable RS/6000 SP system limits for the number of frames.
- (#9222) NODE ATTACHMENT ETHERNET BNC BOOT FEATURE - This feature will get a BNC cable to allow RS/6000 SP system Ethernet communications and booting with your SP-attached server, whether switch-attached or not.
- (#9223) NODE ATTACHMENT ETHERNET TWISTED PAIR BOOT FEATURE - This feature tracks the choice to incorporate the SP-attached servers as part of an RS/6000 SP system Ethernet Twisted Pair network,

but it provides no twisted pair cable. As in the past, the customer is responsible for providing their own twisted pair Ethernet cables.

- (#5765-D51) PSSP 3.1 Software - feature for each SP-attached server license, the features are as follows:
  - (#5800) 4mm tape
  - (#5801) 8mm tape
  - (#5803) CD ROM

  The charge features are different - as follows:
  - (#4001) PSSP software on an RS/6000 SP system
  - (#4002) PSSP software on an Externally Attached Server

Note

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

For full details on configuring an S-Series server as a node in an SP system, refer to *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280.

2.6 Adapter restrictions

The S-Series PCI slot population rules are very complex and cannot easily be explained in this book. Refer to the *PCI Adapter Placement Reference*, SA38-0538, for advice and guidance on the supported configurations of adapters in RS/6000 systems. The placement reference contains information on both maximum configurations and optimal configurations for best performance.

Extensive configuration rules and checking procedures have been incorporated into the PCRS6000 configurator to ensure valid system configurations. Configurations generated without utilizing the PCRS6000 configurator may create orders that cannot be manufactured, resulting in possible order rejection and/or delayed delivery.

There are 14 PCI slots split across 4 PCI buses in each I/O Drawer. Slots assigned to each bus are: 1-4, 5-8, 9-10, and 11-14. Slots 1, 5, 9, 10, and 14

are 64-bit slots. The remaining slots are 32-bit. 32-bit adapters can be used in 64-bit slots.

System maximum limits for adapters and devices may not provide optimal system performance. These limits are given for connectivity and functionality assurance.

Configuration limitations have been established to help ensure appropriate PCI bus loading, adapter addressing, and system and adapter functional characteristics when ordering primary and secondary I/O Drawers. These I/O Drawer limitations are in addition to the individual adapter limitations shown in the feature descriptions section.

The service processor must always be located in slot 8 of the primary I/O Drawer. To ensure proper functionality of the service processor, certain other adapters, such as SSA adapters, cannot share a PCI bus with the service processor.

The service processor provides two serial ports for direct console attachment and support processor use only. These ports are not to be utilized for other functions, such as HACMP heartbeat cables or UPS interface.

Only two SSA adapters are allowed per PCI bus, with a maximum of 26 per system. A maximum of five SSA adapters are allowed in the primary I/O Drawer while up to seven are allowed in each secondary I/O Drawer.

The S/390 ESCON Channel Adapter (#2751) is limited to four per system and two per I/O Drawer. These adapters must be located in drawers attached via the primary I/O loop for addressing considerations.

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

The following adapters are limited to one per system:

- (#2708) Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter
- (#2830) POWER GXT130P Graphics Adapter
- (#2838) POWER GXT120P Graphics Adapter
- (#6322) Service Processor (Model S70 Advanced)
- (#6326) Service Processor (Model S80)
- (#8396) RS/6000 SP System Attachment Adapter

Chapter 3. Hardware architecture

Models in the S-Series have two or more physical enclosures connected by a power controller and I/O cables. The first enclosure is the CEC, which contains processor cards, memory, a memory controller, and a Remote I/O (RIO) controller, along with power supplies and cooling fans. The components in the CEC are packaged between metal sheets to form what are called books. Book packaging provides greater protection for the cards, ensures that they are seated properly, and helps achieve optimum air flow within the CEC.

The second enclosure is a standard 19-inch black rack with a door added for aesthetics. The rack contains the primary SCSI I/O Drawer. Up to four drawers are supported per system. The second I/O Drawer can be housed in the same rack as the first or in an additional rack. No more than two SCSI I/O Drawers are supported in a single rack.

3.1 RS64 III processor and card

The Model S80 uses the 450 MHz RS64 III processor, which is based on the 262 MHz RS64 II used in the Model S70 Advanced. Both are 64-bit, PowerPC-compatible, four-way superscalar implementations. They are optimized for commercial workloads. The Model S80's processors in particular are designed to perform well running applications that place heavy demands on system memory. The RS64 III architecture addresses both the need for very large working sets and low latency. Latency is measured by the number of CPU cycles that elapse before requested data or instructions can be utilized by the processor.

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

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

Using a technique called Double Data Rate (DDR), the new 8 MB Static Random Access Memory (SRAM) used for L2 is capable of transferring data

twice during each clock cycle. The L2 interface is 32 bytes wide and runs at 225 MHz (half processor speed), but, because of the use of DDR, it provides 14.4 GBps of throughput.

Figure 16 contains a diagram of the processor card layout used on the Model S80. There are six RS64 III processors per CPU card sharing two system buses that connect to the memory controller complex.

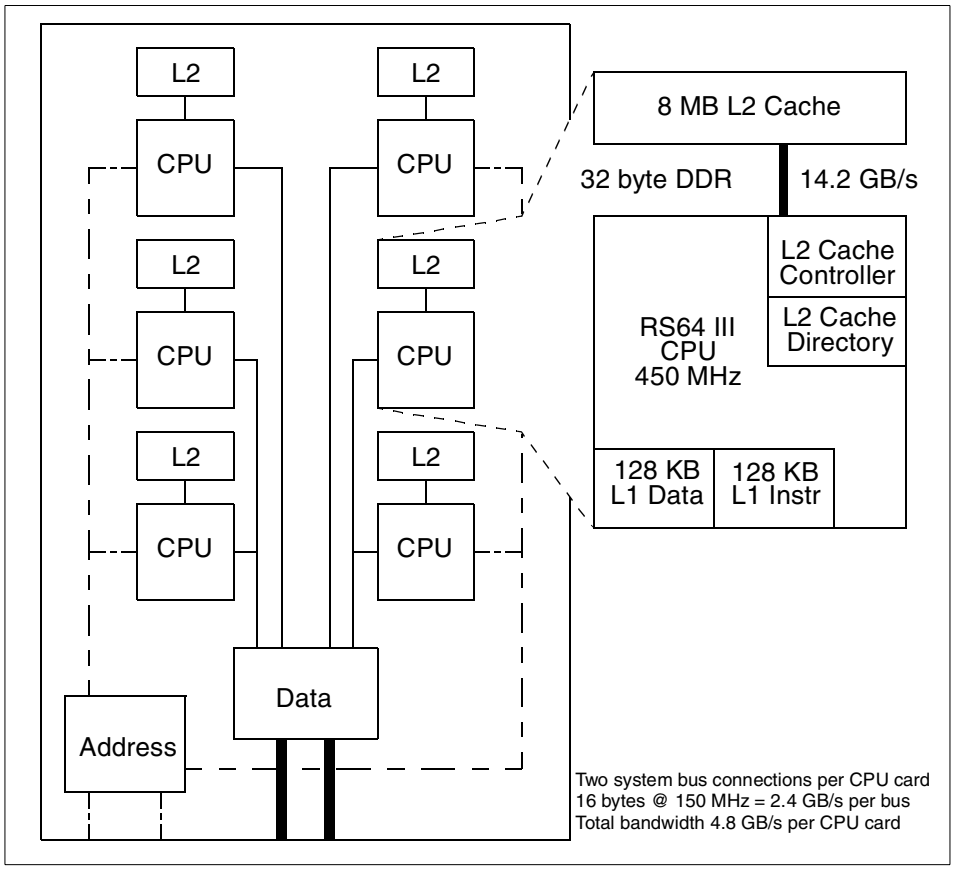

Figure 16. RS64 III processor card

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

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

In summary, the RS64 III features include:

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

3.2 Model S80 memory controller complex

Ten system data buses link the Model S80 processor cards to the memory controller complex and the RIO hub. Each processor card has two buses that connect to two interconnected data flow switches. Each switch connects directly to two quads of Synchronous Dynamic Random Access Memory (SDRAM). Each switch also connects to its own dedicated port on the RIO hub.

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

Figure 17 on page 58 contains a diagram of the memory controller complex layout. The eight system buses connecting the processor cards and the two system buses connecting the RIO ports are 128 bits wide and run at 150 MHz to provide a total bandwidth of 24 GBps. Memory ports are 512 bits wide and operate at 75 MHz. The four memory ports have an aggregate bandwidth of 19.2 GBps. The total memory controller complex switch bandwidth is an impressive 43.2 GBps. Transfer buffers are used in the switch to queue traffic if the needed connections are not immediately available.

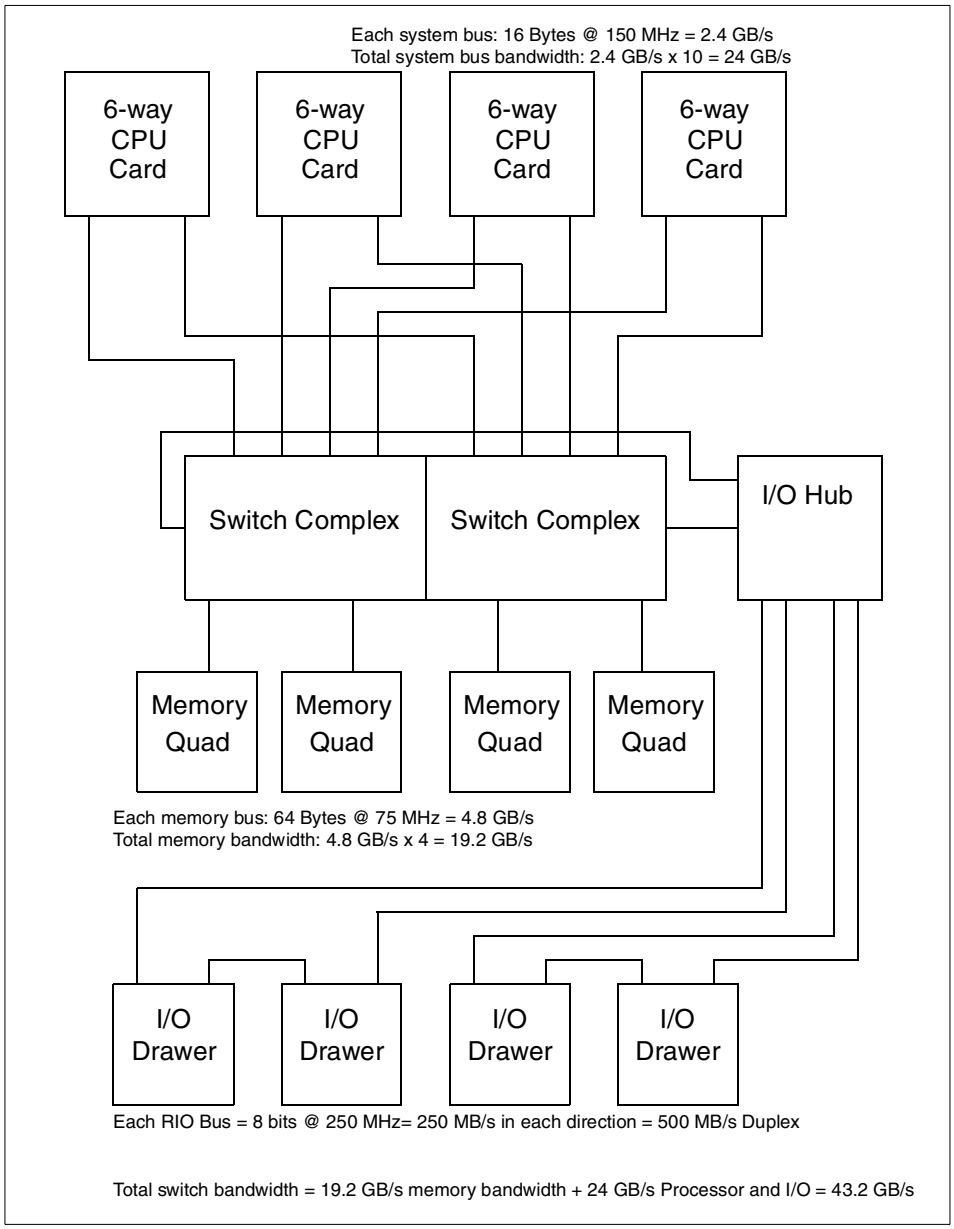

Figure 17. Model S80 system switch complex

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

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

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

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

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

3.3 RS64 II processor and card

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

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

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

The RS64 II features include:

- 64 KB on-chip L1 instruction cache
- 64 KB on-chip L1 data cache with one-cycle load-to-use latency
- 8.4 GBps L2 cache bandwidth
- 32 byte on-chip data buses

- 262 MHz operating frequency
- 4-way superscalar design
- Five stage deep pipeline

Figure 18 contains a diagram of the processor card layout used on the Model S70 Advanced. There are four processors per CPU card sharing a system bus that connects to the memory controller complex.

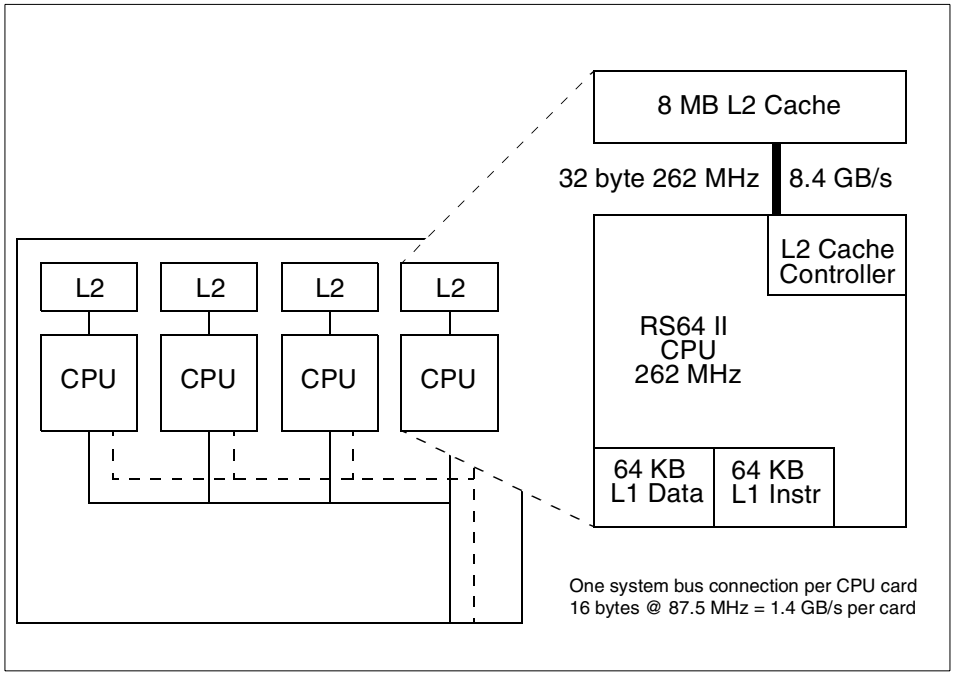

Figure 18. RS64 II processor card

3.4 Model S70 Advanced memory controller complex

Four system data buses link the Model S70 Advanced processor cards to the memory controller complex and the RIO hub. Each processor card has a single system data bus that connects to the data flow switch. The data flow switch connects to the memory subsystem through two memory ports. It also has a connection to the RIO hub. The switch consists of four data flow switch chips, and a separate data flow control chip. The data flow switch chips are the core of the memory controller. An SMP bus arbiter chip in the controller complex prevents switch chip contention.

Figure 19 contains a diagram of the layout of the memory controller complex.

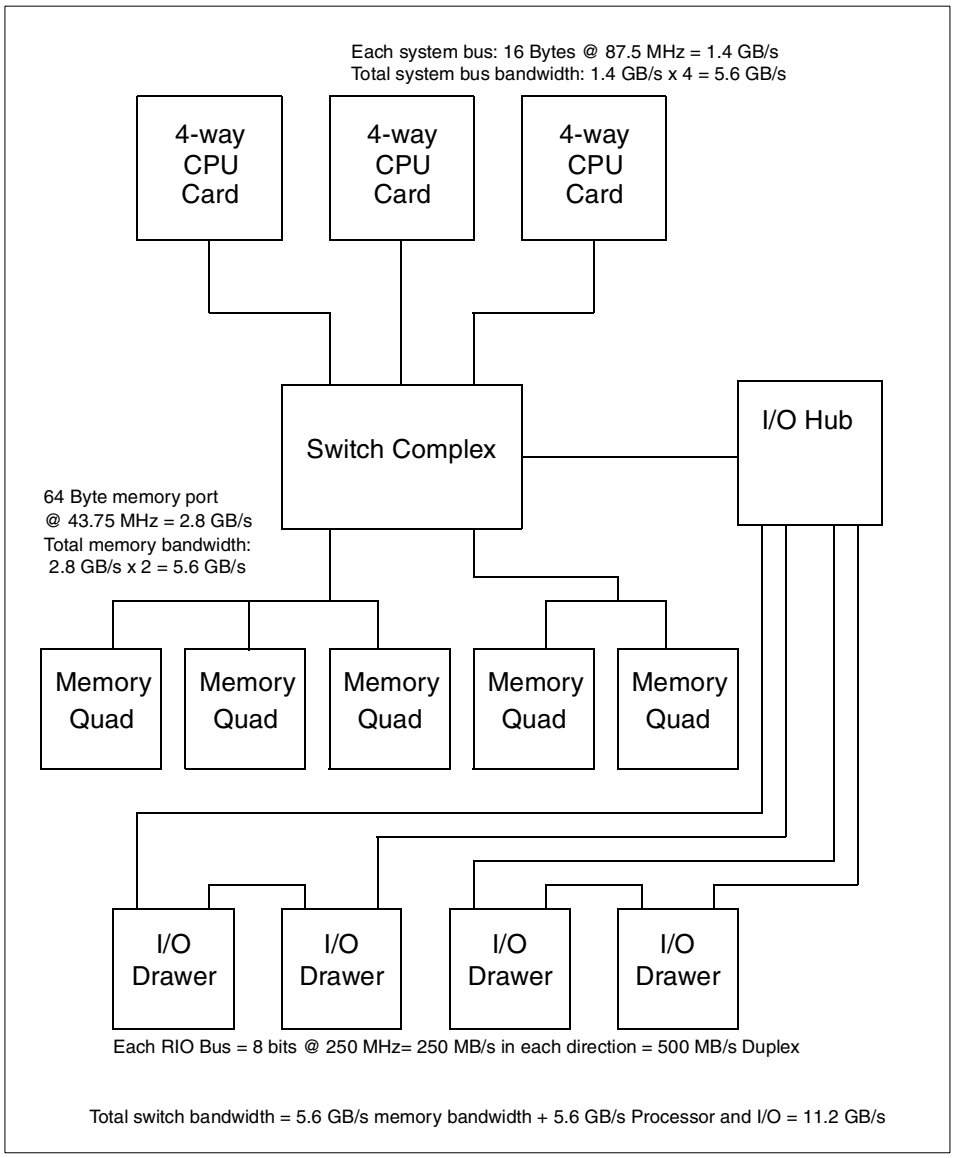

Figure 19. Model S70 Advanced system switch complex

The three system buses connecting the processor cards and the system bus connecting the RIO ports are 128 bits wide and run at 87.5 MHz to provide a total bandwidth of 5.6 GBps. Memory ports are 512 bits wide and operate at

43.5 MHz. The two memory ports have an aggregate bandwidth of 5.6 GBps. The total memory controller complex switch bandwidth is 11.2 GBps.

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

3.5 Memory

System memory for the S-Series comes in books that are plugged into reserved slots in the CEC. Some slots are accessible from the front of the CEC and some from the rear. See Section 2.1.4, “Memory” on page 14, for details of the memory quad layout of the Model S80. See Section 2.2.3, “Memory” on page 24, for details of the memory quad layout of the Model S70 Advanced.

The base Model S80 comes with 2 GB of SDRAM. The base for the Model S70 Advanced is 1 GB. The maximum memory on a Model S80 is 64 GB; on a Model S70 Advanced, it is 32 GB.

The S-Series models use a form of memory designated R1. Memory comes in quads that must be installed in specific locations within the CEC. Memory modules are soldered inside their protective books, which minimizes failures and faults caused by connectors or sockets. Memory from a Model S70 Advanced must be replaced when upgrading to a Model S80 because the latter runs at a faster speed.

The slot layout differs somewhat between the Model S80 and Model S70 Advanced. There are four quads on a Model S80: A, B, C, and D. The Model S70 Advanced has those quads and one more, quad E. The E quad can optionally accommodate much slower DIMM-based memory from older systems being upgraded to a Model S70 Advanced. If DIMM memory is used, quads B, C, and D must be empty, and R1 memory must be installed in quad A. Although this option allows the use of memory from the system being upgraded, it restricts the maximum capacity and performance of the memory subsystem.

On the Model S70 Advanced, the memory system is accessed through two separate ports. Optimization is achieved by equalizing (as much as possible) the amount of memory on each of the ports. Quads A and D are on one port, and quads B, C, and E are on the other. To ensure that both memory ports are used by the system to maximize memory bandwidth, it is important that Model S70 Advanced systems be configured with a minimum of two memory quads.

On the Model S80, memory is accessed through four separate ports. Each half of the system switch complex is connected to two quads. Refer to Figure 17 on page 58 for a graphical representation. Memory quads A and D are on one half of the switch, and quads B and C are on the other half. To ensure that both halves of the memory controller complex are used by the system to utilize the available switch and system bus bandwidth, it is important that Model S80 systems be configured with a minimum of two memory quads.

To get best the performance from either S-Series server, sort all memory features by quad size from largest to smallest. Install the largest memory feature in quad A, insert the next largest memory feature in quad B, and so on. Continue inserting memory quads in the available slots alphabetically, with the smaller features following the larger.

On the Model S80, memory cards M01 through M03 and M08 through M10 are accessible from the front of the system, and M11 through M15 and M20 through M24 are accessible from the rear. Refer to Figure 2 on page 15 and Figure 3 on page 16 for details of the physical layout. Memory for the Model S80 comes in increments of 1 GB (4 x 256 MB), 2 GB (4 x 512 MB), 4 GB (4 x 1024 MB), 8 GB (2048 MB), and 16 GB (4 x 4096 MB). All cards in a quad must be the same size.

On the Model S70 Advanced, memory cards M01 through M06 and M10 through M15 are in the front of the CEC, and M16 through M19 and M21 through M24 are in the rear. Refer to Figure 4 on page 26 and Figure 5 on page 27 for details of the physical layout. R1 memory for the Model S70 Advanced comes in increments of 512 MB (4 x 128 MB), 1 GB (4 x 256 MB), 2 GB (4 x 512 MB), 4 GB (4 x 1024 MB), and 8 GB (4 x 2048 MB).

The memory subsystem has single-bit error correction and double-bit error detection (ECC). The Model S80 memory cards feature redundant modules that will support up to one in 14 memory modules not working. The memory is scrubbed by the controller when it is idle, a feature that is designed to detect soft errors and reduce hard failures.

3.6 Remote I/O connections

On the Model S80, the Remote I/O (RIO) subsystem connects to the CEC via the I/O hub chips on a new replaceable I/O interface card. On the Model S70 Advanced, the chips are mounted directly on the backplane. In either case, four RIO connections are supported allowing a maximum of two RIO loops. The RIO connections are scalable, high-speed, point-to-point interfaces designed for low-latency high-bandwidth connections between two boards or

boxes. Each RIO bus supports up to 500 MB/s total or 250 MB/s in each direction concurrently. RIO cables connect the CEC to the I/O devices located in the I/O Drawers. The RIO connections are set up as loops. The I/O hub chips direct the traffic around the loop in an optimal way for performance, and they redirect traffic if there are link errors.

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

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

3.7 I/O Drawer

The S-Series servers use an enhanced I/O Drawer. The drawers have fully redundant power supplies and fans that can be serviced without shutting down the system. The drawers improve overall cooling by using more powerful, variable-speed fans that increase air flow when one fails. An LED panel on each drawer displays status information.

The primary I/O Drawer for Model S80 and Model S70 Advanced systems contains:

- The I/O planar
- The service processor
- The native I/O card
- Two independent hot-pluggable disk bays (six-packs)
- One available media bay
- One floppy disk drive
- One CD-ROM
- 14 PCI slots (11 available)
- One keyboard port
- One mouse port
- Two serial ports and one parallel port

In the primary I/O Drawer, one PCI slot will be used for the service processor, and two will be used by the controllers for the media bay and first disk six-pack. If the second six-pack is used, a third controller is required. However, one PCI Dual Channel Ultra2 SCSI adapter (#6205) can be used to

replace two PCI Single-Ended Ultra SCSI adapters (#6206), thus, freeing one adapter slot. The dual channel controllers will run at Ultra speed. The 14 PCI I/O slots consist of five 64-bit and nine 32-bit PCI slots. The 64-bit PCI slots are 1, 5, 9, 10, and 14.

The original seven EIA S70 I/O Drawers are not supported on the Model S80.

On the S-Series I/O Drawers, the majority of the function is implemented on the I/O planar. In each drawer, a single RIO-to-I/O bridge bus chip converts the RIO bus to the local mezzanine bus. The mezzanine bus is also called the I/O bridge bus. The I/O bridge bus drives four PCI bridge chips, each with its own local PCI bus containing two or four slots each. Each bus works independently. The RIO-to-I/O bridge chips have “IN” and “OUT” RIO ports to enable redundant and chainable loop connection of RIO devices. The I/O bridge bus runs at 66 MHz and has a 528 MB/s bandwidth.

3.8 I/O bridge bus

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

3.9 PCI buses

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

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

3.10 S80 firsts

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

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

The S-Series is the first RS/6000 to use Predictive Failure Analysis to monitor the error rates on DASD, memory, processors, L1 and L2 caches, and remote I/O. The Model S80 is the first RS/6000 system to allocate one out of every 14 main memory chips for use as a replacement in the event of single-bit errors. Redundancy is also provided for data traffic between the CEC and the I/O Drawer via a looped cable. If the cabling to the I/O Drawer is experiencing errors, the hardware will redirect the I/O traffic through an alternate path on the loop.

Chapter 4. Operating an S-Series server

This chapter describes some of the new features of AIX 4.3.3, which are particularly useful in large commercial environments.

Although many of these features can be used on any RS/6000 server, they are designed to be of most benefit on large SMP systems, such as the S-Series machines.

4.1 AIX 4.3.3

AIX is an integrated UNIX 98 Branded operating environment for POWER- and PowerPC-based workstations, symmetric multiprocessor, and scalable parallel computing systems. This environment enables both the development and execution of computing applications across the RS/6000 product line.

AIX Version 4.3.3 is a third modification release building on the industry-leading strength and stability of AIX 4.3. It is also the first release of AIX to be influenced by Project Monterey. Project Monterey is an initiative led by IBM with SCO, Sequent, and Intel to deliver a single UNIX product line that spans RS/6000 systems, IA-32-, and upcoming IA-64-based systems.

AIX Version 4.3.3 with enhanced 64-bit scalability and functionality provides:

- Significant AIX software scalability enhancements for 24-way SMP systems (RS/6000 S80 Server support).
- AIX Workload Management system with a policy-based method for managing system workload and system resources. See Section 4.2, “Workload management” on page 68.
- AIX exploitation of SecureWay Directory for AIX users and groups.
- Increased scalability, performance, capacity, and capability of e-business with Web serving acceleration, Cisco EtherChannel, and Quality of Service (QoS) administration and support
- TCP/IP enhancements, such as SOCKS V5, gratuitous ARP, and Sendmail, upgrade to Version 8.9.3 with anti-spamming features.
- Improved system availability with support for online Journal File System (JFS) backup and concurrent mirroring and striping. See Section 4.3.1, “Online JFS backup” on page 74 and Section 4.3.2, “Concurrent striping and mirroring” on page 76.
- Enhanced RAS and improved serviceability features for problem determination.

- A port of the Sun Solaris 2.5 NIS+ network information management system in addition to the current NIS support.
- IBM AIX Developer Kit, Java Technology Edition Version 1.1.8.
- Enhanced Ease-of-Use capabilities including additional Web-based System Manager Task Guides and SMIT support.
- X11R6.3, the *Broadway release*, OpenGL enhancements, and graPHIGS enhancements

For full details on the enhancements included in AIX Version 4.3.3, refer to the *AIX Version 4.3 Differences Guide*, SG24-2014.

4.2 Workload management

This section describes the AIX Workload Management (WLM) system, which provides a policy-based method of managing system workload and resources.

4.2.1 Overview

Workload management is designed to give the system administrator more control over how the scheduler and virtual memory manager (VMM) allocate resources to processes. This can be used to prevent different classes of jobs from interfering with each other and to allocate resources based on the requirements of different groups of users.

The major use of workload management is expected to be for large systems with many CPUs and large amounts of memory, such as Model S70 Advanced and Model S80 systems. These large systems are often used for *server consolidation*, where workloads from many different server systems, such as print, database, general user, and transaction processing systems, are combined into a single large system to reduce the cost of system maintenance and administration. These workloads often interfere with each other and have different goals and service agreements. Workload management is designed to address these problems. The same issues can occur in a single environment where the user base has very different system usage characteristics or the system managers have different priority user communities.

Another use of workload management is to provide isolation between user communities with very different system behaviors. This can prevent effective starvation of workloads with certain behaviors, such as interactive or low CPU

usage jobs, from workloads with other behaviors, such as batch or high memory usage jobs.

These usages have two different issues that workload management must address. First, targets for the amount of resources available to different workloads are required. These targets are not absolute; rather, they should be achieved over the long term to provide a degree of fairness. The second issue that workload management must address is boundaries on the amount of resources that a workload can receive. These boundaries can be in terms of maximum resources available and minimum resources that must be made available. These boundaries are not intended to be the major means of separating workloads; rather, they are intended to address special situations where targets are not enough to provide sufficient isolation.

4.2.2 Managing resources with WLM

WLM monitors and regulates the CPU utilization and physical memory consumption of the threads and processes that are active on the system. Using a set of class assignment rules provided by the system administrator, new processes are automatically assigned to a class by WLM upon execution. Classes are grouped into tiers, which indicate the relative importance of each class. Each class has a set of minimum and maximum limits for each resource managed by WLM. In addition, each class has a target value for each resource. This target is representative of the amount of the resource that would be optimal for the jobs in the class.

4.2.2.1 Classes

A class defines common resource requirements for a group of processes. It is possible to create up to 29 classes with up to 16 characters for the class name.

The system administrator defines a set of rules that allow the system to determine which class a process belongs to. Class assignment rules are based on the user ID, group ID, and executable name of the process.

When WLM is enabled, the assignment of jobs to classes according to the class rules is automatic. WLM was designed as a *Set & Forget* Administration.

There are two default classes that WLM creates at initial environment setup:

System Class This class defines the resources for the operating system. It includes all kernel processes and all privileged processes that are not automatically assigned to another class.

Default Class This default class is used as a catch-all at the end of the class rules file. All processes that are not assigned to any other class will be assigned to the default class.

4.2.2.2 Tiers

Tiers are the values that indicate the relative importance of classes to WLM. Each class belongs to a tier, with the importance defined by the tier value, which ranges from 0 to 9 (0 being the most important tier and 9 being the least important).

Tier numbers allow the grouping of classes by equivalent importance. Classes in less important tiers get resources that classes in more important tiers do not require.

4.2.2.3 Shares

The target for usage of different types of resources is specified with shares. The shares are specified as relative amounts of usage between the different classes.

If a class does not have any processes in it, the available resource is divided among the remaining classes according to their share values.

Possible share values range from 1 to 65,000. The default target share, if unspecified, is 1.

Figure 20 shows what happens when a class with 5 shares has no active jobs. The resources that would have been used by the class are made available to other classes.

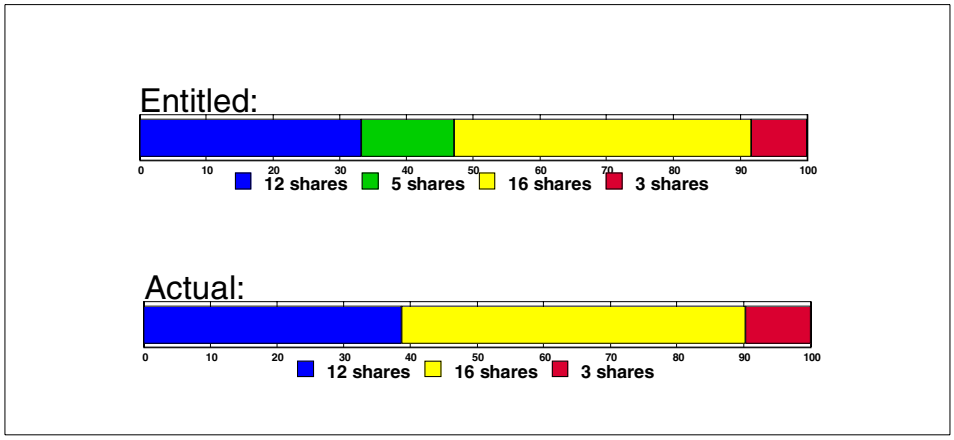

Figure 20. Example share distribution

Each of the active classes still has the same share value, but the real percentage of resource allocated is larger, since there is a smaller total number of shares.

4.2.2.4 Minimum and maximum resource limits

The different resources controlled by WLM can be limited by the following values:

- The minimum percentage of the resource that must be made available when requested. The possible values are integers from 0 to 100. If unspecified, the default value is 0.
- The maximum percentage of a resource that can be made available, even if there is no contention for the resource. The possible values are integers from 1 to 100. If unspecified, the default value is 100.

Keep in mind that the minimum limit reserves resources and the maximum limit restricts resources. The resource target for a class is determined by the number of shares for the class and the number of shares for other active classes. The resource target may also be increased or reduced because limits take precedence over shares.

Resource limit values are specified by resource type in the resource limit file within stanzas for each class. The limits are specified as a minimum to maximum range separated by a hyphen (-) with whitespace ignored. Each limit value is followed by a percent sign (%).

WLM does not place hard constraints on the values of the resource limits. The following are the only constraints:

- The minimum range must be less than or equal to the maximum range.
- The sum of the minimum of all classes within a tier cannot exceed 100.

WLM enforces the maximum range to ensure that a class or process within a class is not given more resource than the specified value. Note that, in the case of a memory constraint, swapping performance can become very poor for processes within the constrained class. Memory minimums for other classes should be used before imposing a memory maximum for any class.

A minimum value constraint on a class means that processes within the class are always allowed to get resources up to the minimum. WLM cannot guarantee that processes actually reach their minimum limit. This depends on how the processes use their resources and on other limits that may be in effect. For example, a class may not be able to reach its minimum CPU entitlement because it cannot get enough memory.

Figure 21 shows a diagram of a class with the limits setup.

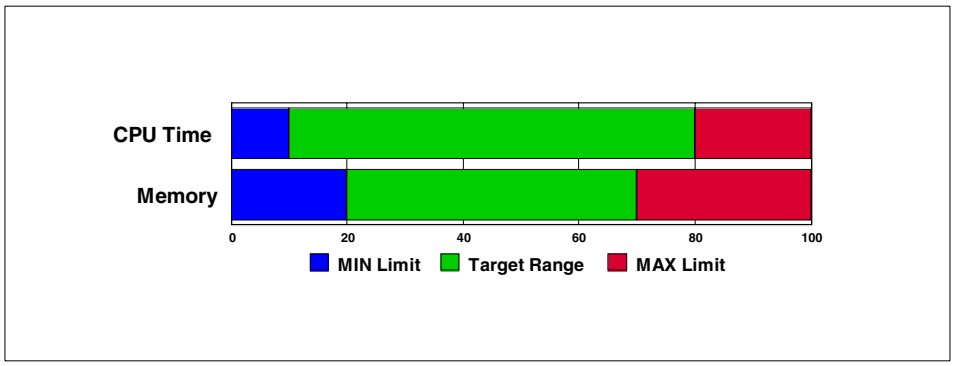

Figure 21. Class with independent resource limits

Each class has independent limits and share values for each resource being controlled by WLM. For example, the class shown in Figure 21 has maximum limits of 80 percent for CPU and 70 percent for memory.

4.2.3 Planning for WLM

WLM offers a fine level of control over resource allocation. However, it is easy to set up conflicting values for the various parameters and obtain undesirable system behaviors. The following tips can help you avoid creating conflicts:

- Know your user base and their basic computing needs when defining classes and class assignment rules.
- Know the resource needs of the main applications.
- Use targets rather than minimum and maximum limits. Targets give the system greater flexibility than hard limits, and targets can help prevent starving applications.
- Try to balance the load using only targets, and monitor the system with the `wlmstat` command. Apply minimum limits for classes that do not receive sufficient share.
- Prioritize some jobs by using tiers.
- Use maximum limits only as a last resort to restrain applications that consume large quantities of resource. Maximum limits can also be used to place hard limits on users' resource consumption (for example, for accounting purposes)

4.2.4 Starting WLM

WLM is an optional service of AIX and can be started manually or automatically. The `wlmcntrl` command allows you to start and stop WLM.

All processes existing in the system before WLM is started remain unclassified and are not monitored by WLM. A future PTF to the WLM system will add the ability to reclassify processes already running when WLM is started.

In normal system mode, it is best to start WLM early in the system initialization process. For example, WLM could be started by an inittab entry, such as the following:

```
wlm:2:once:wlmcntrl -d /etc/wlm/standard > /dev/console 2>&1
```

4.2.5 WLM user interfaces

The administration of WLM can be performed using Web-based System Manager (wsm), SMIT, the command line, or by modifying configuration files directly.

This provides the administrator flexibility to include WLM commands in shell scripts or use a graphical interface to set up or modify the WLM configuration.

Figure 22 on page 74 shows the WLM control application using the wsm user interface.

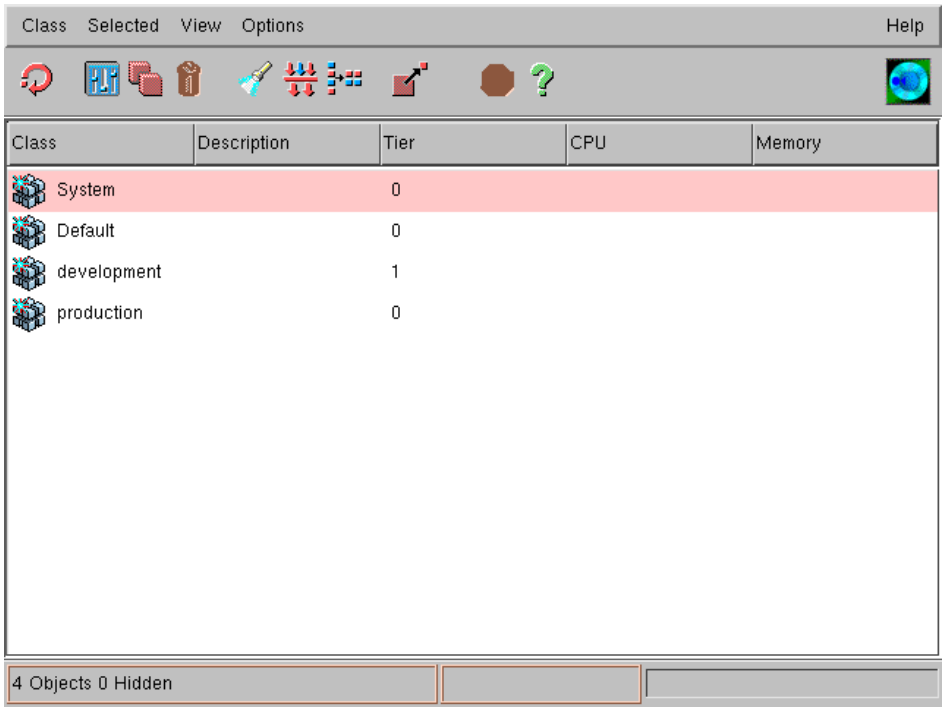

Figure 22. WLM graphical interface

4.3 Additional AIX features

This section describes some of the new AIX features that will be of great benefit in the large commercial server environments commonly found on Model S70 Advanced and Model S80 machines.

4.3.1 Online JFS backup

The Journaled File System (JFS) has been enhanced to support file system online backup. This capability allows a mirrored copy of a file system to be used for backup purposes. A mirror copy of the file system is split off, mounted read-only, and available for backup. The primary copy of the file system is still mounted and in use. This enables the system administrator to back up a consistent copy of the file system data while another copy is still mounted and in use by users and applications. After the backup is complete, the administrator can reintegrate the backup mirror copy and resynchronize it with the other mirror copies.

There are some conditions that have to be met in order to perform an online backup.

- The JFS logical volume and JFS log logical volume must be mirrored.
- The number of mirrored copies of both the JFS logical volume and the JFS log logical volume must be the same.

Note

Splitting a mirrored copy of a file system means that one copy is temporarily dedicated to backup activities and is not available to provide data availability. It is recommended that you have a triple mirror of the file system so you can recover from any disk problem that happens before the backup copy has been reintegrated with the file system.

It is recommended to keep file system activity as minimal as possible during JFS online JFS backup events.

After the backup activities are finished, it is possible to remove the read-only file system in order to have maximum data availability.

The process of splitting off and mounting the backup file system is accomplished by use of the `chfs` command. The correct set of options to this command will split off the required copy of the mirrored logical volume and mount the copy in read only mode at a specified mount point. For example, assuming that /data is a file system on a triple mirrored logical volume, the following command splits off the 3rd copy of the mirror and mounts it read only as /backup:

```
# chfs -a splitcopy=/backup -a copy=3 /data
```

When the `chfs` command completes, the snapshot copy of the file system is mounted and ready for the backup procedure to commence.

Once the backup procedure has completed, the snapshot copy of the file system needs to be reintegrated with the live data. This is performed in two steps. The first step is to unmount the snapshot copy of the file system. The second step is to use the `rmfs` command to remove access to the snapshot copy. At this point, the snapshot data is not actually deleted. The `rmfs` command removes the temporary logical volume device, which was created to access the snapshot copy. It then starts the syncvg process in the background to resynchronize the snapshot copy with the live data. For example:

```
# umount /backup
# rmfs /backup
```

Only the data blocks that have changed since the snapshot was taken are actually synchronized. The duration of this operation depends on a number of factors including the size of the file system, the number of blocks that have changed, and the speed of the disk subsystem on which the file systems are stored.

4.3.2 Concurrent striping and mirroring

For some time, the Logical Volume Manager (LVM) in AIX has provided the ability to create striped or mirrored logical volumes, although a logical volume could not be both striped and mirrored. In this AIX release, the LVM combines RAID 1 (mirror) data availability with RAID 0 (striped) performance by supporting (entirely in software) a striped logical volume with mirrors. This feature further enhances data availability in high-performance striped logical volumes by tolerating disk failures. The remaining disks in the striped mirror copy continue to service striped units contained on these disks. The replacement of a disk where only the partitions on the new disk are synchronized is provided through the `migratepv` and `replacepv` commands.

In addition, all logical volumes can now utilize a new partition allocation policy called *Super Strict*. This Super Strict policy does not allow partitions from one mirror to share a disk with any partitions from a second or third mirror, thus, helping to further reduce the probability of data loss resulting from a disk failure.

These new functionalities are not backward compatible; therefore, new volume groups supporting these features cannot be imported and used with previous versions of AIX.

4.3.2.1 RAID Levels 0 and 1

RAID 0 is also known as data striping. Conventionally, a file is written out sequentially to a single disk. With striping, the information is split into chunks and the chunks written to (or read from) a series of disks in parallel. There are two main performance advantages to this:

- Data transfer rates are higher for sequential operations due to the overlapping of multiple I/O streams.
- Random access throughput is higher because access pattern skew is eliminated due to the distribution of the data. This means that with data distributed evenly across a number of disks, random accesses will most likely find the required information spread across multiple disks and, thus, benefit from the increased throughput of more than one drive.

Although RAID configurations are normally associated with availability, keep in mind that RAID 0 is only designed to increase performance.

RAID 1 is also known as disk mirroring. In this implementation, duplicate copies of each chunk of data are kept on separate disks. If any disk in the array fails, the mirrored twin can take over. Read performance can be enhanced as the disk with its actuator closest to the required data is always used, thereby, minimizing seek times.

The response time for writes can be somewhat slower than for a single disk depending on the write policy; the writes can either be executed in parallel for speed or serially for safety. Writing in parallel means that the write process will complete in the time taken for the slowest drive to finish in the mirrored pair; this is quick, but it means that a failure in writing to one of the pair of disks is not immediately detectable. In contrast, writing sequentially means that the write to the mirrored copy is not initiated until the first write has successfully completed. This is slower, but any errors are immediately detectable. Mirroring improves response time for read-mostly applications and improves availability at the expense of cost, since twice (or three times) as many disks are required as disk space.

RAID 1 is best-suited to applications that require high data availability and good read response times and in settings where cost is secondary.

Table 11 compares the relative characteristics of different RAID levels.

Table 11. Characteristics of different RAID levels

RAID Level	Availability mechanism	Capacity	Performance	Cost
0	none	100%	high	medium
1	mirroring	50%	medium	high
0+1	mirroring	50%	high	high
5	parity	80%	medium	medium

It is important to keep in mind that each different level of RAID configuration has advantages and disadvantages. A combination of striping and mirroring provides high performance and high availability, but there is a high cost.

Note

RAID 0 mainly improves sequential workloads. For heavy random workloads and high availability needs, you should consider using RAID 1 only.

4.3.3 Fast device configuration

The AIX configuration manager now allows multiple device configuration methods to run in parallel during system boot. This will produce a faster reboot when multiple devices of the same type, such as multiple SCSI disks, TTYs, and multiport asynchronous adapters, are connected to an AIX system. A serialization mechanism is used when the configuration manager recognizes that new devices have been added to the system. This causes the new devices to be sequentially configured.

Systems with a large number of asynchronous I/O adapters benefit from a reduction in both machine boot time and duration of the `cfgmgr` command. Generally, each asynchronous I/O adapter can take up to five minutes to be configured. With this new feature, up to 16 device configuration methods can run in parallel.

4.4 General AIX improvements for large systems

Some AIX improvements will be particularly beneficial for large SMP systems. This section gives a brief description of some of these changes.

4.4.1 Increased threads and process limits

In previous AIX versions, the number of system threads and processes were each limited to 131072. In some large SMP systems, the number of threads and processes could reach the limit. In AIX 4.3.3, these numbers have been extended to 524288 for threads and 174080 for process limits.

4.4.2 Multiple run queues with load balancing

On SMP systems, AIX 4.3.3 implements multiple run queues. Each processor now has a local run queue in addition to the original system-wide global run queue.

When new threads are created, they are added to the system-wide global run queue. When a processor is looking for a thread to dispatch, it will examine the global run queue in addition to its local run queue and choose the highest priority thread. Once a thread is on a processors local run queue, it will tend

to stay on that run queue and, thus, always be run on the same processor. This increases the probability that the data and instructions required to run the thread will still be in the cache of the processor, thus, reducing the need to fetch instructions and data from main memory.

Periodic load balancing takes place automatically to ensure that all processors have similar numbers of threads in their local run queues and to prevent threads from bottle necking on a single processor.

The multiple run queues improve system performance by dramatically reducing lock contention for the global run queue and increasing processor affinity for individual threads. In turn, this reduces the overhead on the cache and memory subsystem, thus, allowing greater system throughput.

4.4.3 SMP scalability improvements

In addition to the changes to the run queue system, many other changes have been made to the process management system. Most of the changes are aimed at minimizing lock contention in the kernel, which is of much more importance now that 24-way SMP systems are supported. Some of the changes include analyzing and improving the code being run while particular kernel locks are held. Other situations called for implementing new more granular locks to reduce contention. One area to benefit from this is the process table lock.

4.5 PCI Dual Channel Ultra2 SCSI Adapter (#6205)

The PCI Dual Channel Ultra2 SCSI Adapter is a 64-bit adapter and is an excellent solution for high-performance SCSI applications. The PCI Dual Channel Ultra2 SCSI Adapter provides two SCSI channels (buses). Each SCSI bus can either be internal or external and supports a data rate of up to 80 MBps, up to twice the maximum data transfer rate of previous Ultra SCSI adapters, which was 40 MBps.

In order to achieve an Ultra2 SCSI bus data rate of up to 80 MBps and also maintain a reasonable drive distance, the adapter utilizes Low Voltage Differential (LVD) drivers and receivers. In order to utilize this Ultra2 80 MBps performance, all attaching devices or subsystems must also be Ultra2 LVD devices. If any device is not Ultra2 LVD, the adapter will switch its SCSI bus to single-ended (SE) performance and interface at the lower SE SCSI bus data rate of the device.

Two industry-standard VHDCI 68 pin connectors are mounted on the adapter's end bracket allowing attachment of various LVD and SE external

subsystems. The 0.3 meter, VHDCI to P, Mini-68 pin to 68 pin (#2118) converter cable can be used with older external SE subsystems to allow connection to the VHDCI connector on the PCI Dual Channel Ultra2 SCSI Adapter.

Any supported RS/6000 system can be set up to boot from the PCI Dual Channel Ultra2 SCSI Adapter (#6205). If you are running with AIX 4.3.3 or later software, this adapter has native boot support as part of that level of AIX software. If you are running with 4.2.1 software, the following procedure applies in order to boot using the PCI Dual Channel Ultra2 SCSI Adapter:

- The designated boot SCSI disk can be located under the covers of a processor unit or in an external SCSI storage unit.
- AIX version 4.2.1 must be loaded to the designated SCSI boot disk using the AIX Network Install Manager (NIM) before booting from the SCSI boot disk.
- The system with a designated SCSI boot disk must have a network connection with another RS/6000 system performing the NIM Master function to perform the install. On RS/6000 SP systems, a similar network install is performed from a control workstation.
- Once AIX 4.2.1 with updates is installed on the designated SCSI boot disk and the system is configured for booting, booting takes place from the boot disk drive without any support from the control processor or NIM Master, and the system does not have to be connected to the network at boot time.

Chapter 5. Reliability, availability, and serviceability

Reliability, availability, and serviceability (described as a collective term, RAS) describe how well a system can perform its intended function on demand and how quickly problems and errors can be repaired. Commercial servers running mission-critical applications need strong RAS capabilities. RAS encompasses techniques for reducing faults and minimizing their impacts, shortening repair time, and enabling faster problem resolution.

In the S-Series, RAS begins with the development of architectures that give high priority to these characteristics. The emphasis on RAS begins in product development, where designs are tested, evaluated, and optimized. Concern for reliability continues through the manufacturing and distribution processes, where quality is continually evaluated and carefully measured against documented standards. Finally, the focus on RAS extends to service and support, where real-world experience is measured against design criteria. Hardware warranty and maintenance support receive undivided attention within IBM, and significant customer problems are addressed by teams of experts.

The processes IBM follows in designing, testing, manufacturing, and servicing a product are periodically audited and certified for compliance with International Standards Organization (ISO) 9000 guidelines.

5.1 Designed for higher RAS

During the development of the S-Series, a major effort has been made to analyze single points of failure within the central electronic complex. The processors, L1 and L2 cache, system memory, memory controller complex, and the remote I/O subsystem have been designed to provide mainframe-like levels of reliability. They undergo additional stress testing and screening by suppliers above and beyond what is required of industry-standard components used in many UNIX systems today.

The S-Series design aids in the recognition of intermittent errors that are either corrected dynamically or reported for further isolation and repair. Parity checking on the system bus, cyclic redundancy checking (CRC) on the remote I/O bus, and the use of error correcting code on memory and processors contribute to outstanding RAS characteristics. Redundant arrays with cyclic redundancy checking are implemented in the power controller card and the operator panel. The service processor arrays are redundant and use CRC checking.

The Model S80 does memory scrubbing, which is the process of reading the contents of memory during idle time and checking and correcting any single-bit errors that may have accumulated. The Model S80 also does bit steering (replacing failing storage with spares). In the new system, there is a 14:1 ratio between live and spare memory bits.

During the boot sequence, built-in self test (BIST) and power-on self test (POST) routines check the processors, cache, and associated hardware required for a successful system start. These tests run every time the system is powered on.

Additional testing can be selected at boot time to fully verify the system memory and check the chip interconnect wiring. When a system reboots after a hard failure, it performs extended mode tests to verify that everything is working properly and that nothing was compromised by the failure. This behavior can be overridden by the systems administrator.

The S-Series RAS design enhancements include:

- Automatic error capture and problem isolation
- Dynamic error recovery
- Single-bit error correction, double-bit error detection on internal processor arrays, and L1 and L2 cache and system memory
- Bit steering and scrubbing on memory
- Redundant power supplies and cooling fans providing fault tolerance and concurrent maintenance for those subsystems
- Predictive failure analysis on processors, memory, RIO, and disk
- Processor boot-time deallocation of resources based on run-time errors
- Highly-reliable, stress-tested components
- Concurrent diagnostics

5.2 Service processor

The service processor and IBM Service Director for RS/6000 work independently and together to achieve high levels of reliability, availability, and serviceability on the S-Series servers.

An independent microprocessor, the service processor occupies a 32-bit PCI slot in the primary I/O Drawer in the Model S80 and Model S70 Advanced. It comes installed in slot 8. The service processor runs its own firmware and has access to non-volatile memory and hardware components within the

central electronic complex. It can control system behavior under predetermined conditions. The firmware has an asynchronous menu-driven interface that can be accessed either locally or via modem from a remote support site. If desired, the local systems administrator can mirror the remote terminal and keyboard maintaining oversight and eliminating the need to share passwords with service personnel. This is called console mirroring.

During initialization, the service processor monitors diagnostic routines that eventually culminate in control of the system being turned over to the operating system. Even then, the service processor continues to monitor hardware, environmental conditions, and, if so configured, the operating system. If a fatal hardware or software problem brings the system down, the service processor can automatically call support, the customer's own help desk, or a pager number.

During its run-time monitoring of the hardware, if the service processor detects a non-fatal error that could develop into one that would bring the system down, it logs the problem in NVRAM and notifies AIX that it has error information for the error log. On the AIX side, the NVRAM will be accessed by the errdemon, and the problem will be entered in the AIX error log. An error logged in this manner will trigger an error log analysis, and, if Service Director has been configured, a service call will be placed to IBM, and the customer's help desk will be notified of the situation.

If a fatal error occurs during run-time, the system will halt. Depending on whether a reboot policy has been set, the system will attempt to restart itself. The service processor can be configured to call home and report the problem before attempting the restart. Typically, such calls would go to IBM, the customer's help desk, a pager carried by an on-call technician, or some combination of contacts. If the system is restarted, firmware will cause the service processor to deallocate the failing module.

Other capabilities of the service processor include the ability to remotely power-off/power-on the system, read its own and system POST error logs, and read vital product data (VPD). The service processor can change the bootlist, thus, enabling an alternate boot source. It can also view the boot sequence history, which lists the progress indicators that appeared in the operator panel LCD window during the last boot. The service processor also enables the systems administrator to establish password protection of the firmware menus, view the boot sequence history, and change operating system surveillance and reboot policies.

Updates to service processor and system firmware should be performed by trained personnel. The updates are available from the following RS/6000 support page on the Internet:

`http://www.rs6000.ibm.com/support/micro`

Code can be downloaded directly to a server and then installed, or it may be downloaded as a file on a separate workstation. The file is then copied to the Model S80 or Model S70 Advanced and installed. The firmware can be updated remotely using diagnostic Service Aids or AIX command line options.

By default, a backup copy of the prior version of firmware is maintained on the system so that it can be booted with the older version if necessary. At the system administrator's discretion, the new firmware can be copied to the alternate storage location so that the two copies are identical, providing redundant firmware capability in case of corruption.

5.3 Service Director for RS/6000

Service Director ships automatically on all RS/6000s including the S-Series models. There is no charge for its use as long as the system is covered by warranty or an IBM maintenance agreement. Groups of machines can run Service Director with one machine serving as a single point of control and the link to IBM service.

The mission of Service Director is to monitor hardware and analyze recoverable faults and report them. Fatal errors that cause the system to crash can be reported via the service processor. If desired, Service Director can be set up to automatically place a service call to IBM when it detects a problem requiring an on-site visit. Alternatively, it can pass an alert to a help desk so that the customer can decide whether to place a service call, or it could do both. Service Director can send alerts as electronic mail to a limited number of addresses specified by the systems administrator.

Service Director is menu-driven. The screens can be accessed through the Systems Management Interface Tool (SMIT). The program enables more effective management of hardware support, allowing administrators to quickly and easily view the details of problems that have occurred, the status of open calls, and the machine's service history.

Service Director requires a modem. It may share a modem used by the service processor.

Service Director is designed to automatically report hardware problems based on default settings. However, the customer may modify the settings to prevent Service Director from placing a call, such as during hardware upgrades and testing or if the failing component is not covered by an IBM service agreement. The program will also automatically notify an administrator by E-mail before expiration of a warranty or service agreement, giving the customer time to sign or renew a maintenance contract.

5.4 Additional system RAS features

Many server components, such as disk drives, power supplies, and cooling fans, are now hot-swappable; so, many common repairs can be made without stopping applications and taking the system offline. For example, with a RAID controller and hot-pluggable disks, normal operation can continue, perhaps in a slightly degraded mode, while a failed drive is repaired.

The Model S80 brings new levels of availability and serviceability to the enterprise server. Within the CEC, the extra power supply provides redundancy in case of failures in either the bulk or regulated power subsystems. Concurrent repair is supported on the bulk supplies while the regulators will require a controlled system shutdown before repairs can be made. Optional uninterruptible power supply (UPS) systems are supported and recommended for mission-critical servers.

The CEC cooling subsystem has what is called N+1 redundancy. That is, there is one blower more than the number required to keep the system running. If one fails, repairs can be made without shutting down the system. The fans can also adjust their speed to partially compensate for a single failure.

Constant power monitoring hardware assists in the detection of early power loss and notifies the operating system to attempt an orderly shutdown. This same power monitoring hardware detects the loss of redundant bulk power supplies, regulators, fans, and blowers and reports them to the operating system for logging in the system error log and for notification for deferred maintenance.

With enhanced AIX diagnostics available with AIX 4.3.3, system administrators can keep track of diagnostic activity more effectively by using the new diagnostic event log. This log can be viewed by using the Display Previous Diagnostic Results task under concurrent diagnostics. Tasks and service aids within diagnostics have been ordered alphabetically for improved ease of use.

A diagnostic exerciser has been added for processors to enhance problem determination. It provides a means of verifying both memory and processor repairs previously detected by error log analysis.

Auto-restart options for the server, when enabled, are designed to reboot the system automatically following:

- Unrecoverable software error
- Software hang
- Fatal hardware error
- Environmentally-induced (AC power) failure

AIX 4.3.3 removes the prohibition against mirroring a dump device. AIX will not complain about the mirrored devices, but it will only write to and read from the primary copy. Additionally, the Logical Volume Manager supports online backup of journaled file systems.

The latest release of AIX also supports mirroring of striped file systems, enabling RAID 0 + 1 to improve performance and protect against disk failures.

The Kernel Debugger (kdb) is a tool being added to the AIX system to provide a symbolic debugger for the AIX kernel, kernel extensions, and device drivers. There is also a `kdb` command, which is an alternative to the current `crash` command, to allow examination of system crash dumps.

These standard high-availability features, coupled with the optional High Availability Cluster Multi-Processing (HACMP) program, provide unmatched availability in the UNIX marketplace.

5.5 AIX RAS

As part of basic program design, operating systems developed by IBM have historically mandated recovery processing in both the mainline program and separate recovery routines. However, a significant percentage of AIX code originated outside of IBM. Consequently, recovery processing, such as that employed in IBM mainframe operating systems, is not the norm in AIX.

The classical approach to recovery is largely based on the premise of keeping the system or application alive at any cost. This approach grew out of an environment in which each system was essentially stand-alone, supported hundreds of users, and took a long time to restart. Today, even in large multiuser environments, many customers require that recovery processing be subject to a time limit. They have concluded that rapid termination with quick

restart or takeover by another application or system is preferable to a delayed restart of the one that failed.

Serviceability is the ability to diagnose and correct or recover from an error when it occurs. The serviceability capabilities and enablers in AIX are referred to as the software service aids. The primary software service aids are error logging, system dump formatting, and tracing.

The software service aids provide a common set of tools for performing problem determination and problem source identification for AIX software. Problem determination is the name IBM gives to those activities performed in determining whether a malfunction is caused by hardware or software. Those activities that take place after problem determination to determine the specific malfunctioning hardware or software component are referred to as problem source identification.

With proper instrumentation in the software, the software service aids enable the customer to determine whether a problem is caused by a defect in hardware or software. If a problem is due to malfunctioning hardware or software, the service aids lead the customer to place a service call and enable him or her to report an accurate description of the problem. If a problem is not due to malfunctioning hardware or software, the service aids should lead the customer to problem resolution without requiring an unnecessary service call.

Appendix A. Installation requirements

This appendix details the environmental, electrical, and installation requirements of the S-Series systems including physical dimensions, clearances, and power requirements.

A fully-configured S-Series system consists of a CEC, at least one I/O Rack, and at least one I/O Drawer.

A.1 Central Electronic Complex

Table 12 details the specifications and requirements for the Model S80 Central Electronic Complex (CEC). Table 13 on page 90 details the specifications and requirements for the Model S70 Advanced CEC.

The CEC component of Model S80 and Model S70 Advanced systems is only available in 240 Volt single phase configurations. No three-phase options are available. Each CEC requires a single electrical connection.

Table 12. Model S80 Central Electronic Complex

Dimensions	**Metric**	**Imperial**
Height	1577 mm	62.0 in
Width	567 mm	22.3 in
Depth	1041 mm	40.9 in
Weight	**Metric**	**Imperial**
Minimum	400 kg	880 lbs
Electrical		
Power source loading (Maximum in KVA)	2.129 KVA	
Voltage range (V ac)	200 - 240	
Frequency (Hertz)	50 - 60	
Thermal output (Maximum)	6904 BTU/hr	
Power requirements (Maximum)	2023 watts	
Power factor	0.92 to 0.98	
Inrush current[1]	43 amps	

Maximum altitude	2135 m (7100 ft)	
	Operating	**Non-operating**
Temperature Range	10 to 37.8°C	1 to 60°C
	50 to 100°F	34 to 140°F
Humidity (Noncondensing)	8 to 80%	8 to 80%
Wet Bulb Requirements	23°C (73°F)	23°C (73°F)
Noise Emissions	**Operating**	**Idle**
L_{WAd}	7.0 bels	7.0 bels
L_{pAm}	N/A	N/A
$<L_{pA}>_m$	N/A	N/A
Impulsive or prominent discrete tones	No	No

[1] Inrush currents occur only at initial application of power, no inrush occurs during normal power off-on cycle.

Table 13. Model S70 Advanced Central Electronic Complex

Dimensions	**Metric**	**Imperial**
Height	1577 mm	62.0 in
Width	567 mm	22.3 in
Depth	1041 mm	40.9 in
Weight	**Metric**	**Imperial**
Minimum	400 kg	880 lbs
Electrical		
Power source loading (Maximum in KVA)	1.887 KVA	
Voltage range (V ac)	200 - 240	
Frequency (Hertz)	50 - 60	
Thermal output (Maximum)	5796 BTU/hr	
Power requirements (Maximum)	1698 watts	
Power factor	0.9	

Inrush current[1]	102 amps	
Maximum altitude	2135 m (7100 ft)	
	Operating	**Non-operating**
Temperature Range	10 to 37.8°C	1 to 60°C
	50 to 100°F	34 to 140°F
Humidity (Noncondensing)	8 to 80%	8 to 80%
Wet Bulb Requirements	23°C (73°F)	23°C (73°F)
Noise Emissions	**Operating**	**Idle**
L_{WAd}	7.0 bels	7.0 bels
L_{pAm}	N/A	N/A
$<L_{pA}>_m$	N/A	N/A
Impulsive or prominent discrete tones	No	No

[1] Inrush currents occur only at the initial application of power; no inrush occurs during the normal power off-on cycle.

A.2 I/O Rack

Depending on the configuration of the system, you will require at least one S-Series I/O Rack. The specifications of the I/O Rack are the same for the Model S80 and Model S70 Advanced systems. Table 14 shows the specifications for the S-Series I/O Rack.

Each I/O Rack requires an electrical connection for each Power Distribution Unit (PDU). If you configure an I/O Rack with two PDUs, it will require two electrical connections.

Table 14. S-Series I/O Rack

Dimensions	Metric	Imperial
Height	1577 mm	62.0 in
Width	650 mm	25.5 in
Depth	1019 mm	40.1 in
Weight	**Metric**	**Imperial**

Base rack	159 kg	349 lbs
Electrical	See specifications for drawers or enclosures	
Temperature Range		
Humidity Requirements		
Noise Emissions		

A.3 10 EIA I/O Drawer

Model S80 and Model S70 Advanced systems use an identical 10 EIA I/O Drawer. The I/O Drawer has dual redundant power supplies and requires two connections on the PDU of the I/O Rack in which it is installed. Table 15 details the specifications and requirements of the I/O Drawer.

Table 15. S-Series 10 EIA SCSI I/O Drawer

Dimensions	**Metric**	**Imperial**
Height	440 mm	17.3 in
Width	443.2 mm	17.5 in
Depth	843.2 mm	33.2 in
Weight	**Metric**	**Imperial**
Minimum configuration	89 kg	195 lbs
Maximum configuration	93 kg	206 lbs
Electrical		
Typical power source loading (in KVA)	0.4 KVA	
Maximum power source loading (in KVA)	1.0 KVA	
Voltage range (V ac)	200 - 240 (autoranging)	
Frequency (Hertz)	50 / 60	
Thermal output (Typical)	1228 BTU/hr	
Thermal output (Maximum)	3071 BTU/hr	
Power requirements (Typical)	360 watts	
Power requirements (Maximum)	900 watts	

Power factor	0.9	
Inrush current	170 amps	
Maximum altitude	2135 m (7000 ft)	
	Operating	**Non-operating**
Temperature Range	10 to 40°C	10 to 52°C
	50 to 104°F	50 to 125.6°F
Humidity (Noncondensing) without tape drive	8 to 80%	8 to 80%
Humidity (Noncondensing) with tape drive	20 to 80%	20 to 80%
Wet Bulb without tape drive	27°C (80°F)	27°C (80°F)
Wet Bulb with tape drive	23°C (73°F)	27°C (80°F)
Noise Emissions	**Operating**	**Idle**
L_{WAd}	5.9 bels	5.3 bels
L_{pAm}	N/A	N/A
$<L_{pA}>_m$	N/A	N/A
Impulsive or prominent discrete tones	No	No

A.4 Physical space requirements

Multiple I/O Racks can be placed adjacent to one another, with a 75 to 125 mm (3 to 5 in.) gap. Only the racks at each end of the installation require the full 915 mm (36 in.) clearance. Figure 23 on page 94 details the minimum installation and service clearances required for S-Series servers.

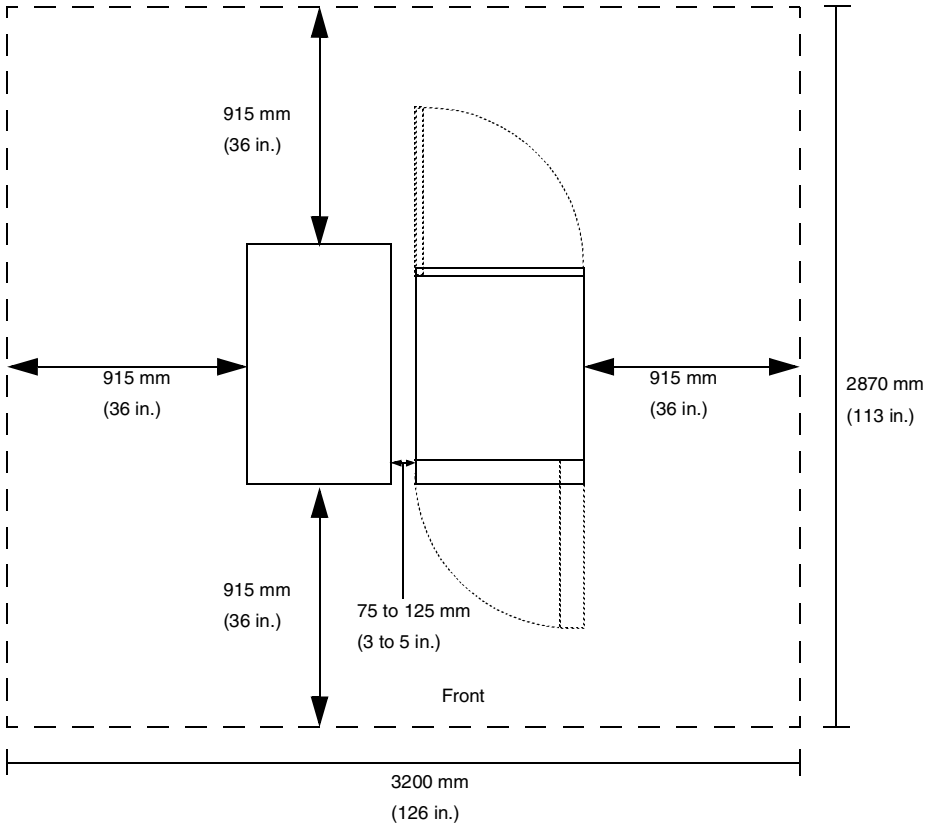

Figure 23. S-Series system service clearances

Appendix B. A practical guide to the service processor

Each S-Series server comes with a service processor to help maintain high availability and reliability and serviceability (RAS). The service processor also makes it possible to perform many maintenance and support tasks from a remote location via modem.

Even when dial-in support is not contemplated, customers should be aware of the capabilities for calling out to report serious system errors. In addition, there are settings in the service processor menus that control the behavior of the system after a power outage and following a hardware or software error. Finally, customers should understand the purpose of system and service processor firmware, and they should know how to get updates.

B.1 How to access the service processor menus

To access the service processor menus, the systems administrator needs to access one of the native serial ports: S1 or S2. If remote dial-in support is enabled, one of the ports needs to have an asynchronous modem attached and configured. The other port may have an ASCII terminal attached, or it may have a leased-line or dial-up connection to a remote customer support location. The service processor menus give an administrator a great deal of flexibility in deciding which ports to enable and whether to allow dial-out only or to support dial-in as well.

B.1.1 Local access

To access the service processor menus from a terminal attached to the S1 or S2 ports, perform the following steps:

1. If the system is running, use the `shutdown` command to bring it to a halt.
2. After the system has powered off, restart it using the white power button on the operator panel.
3. Watch the checkpoints that appear in the operator panel display.
4. Immediately after E04F and just as it enters checkpoint E07A, the system will beep three times. Press any key on the ASCII terminal.
5. Depending on whether a password has been set and which password has been entered, either the Main Menu (privileged access password) or General User Menu (general access password) will appear.

B.1.2 Remote access

To access service processor menus from a remote location, an asynchronous modem must be attached to S1 or S2, dial-in must be enabled on the port, and a TTY must be configured in AIX. Perform the following steps:

1. Dial into the system and use the `shutdown` command to power it off.
2. Use the Ring Indicate Power On (RIPO) utility to restart the system. Dial the telephone number of the modem attached to the serial port and let it ring one to three times before hanging up. By default, the RIPO is set to one ring. The number of rings is alterable through Service Aids in diagnostics.
3. Wait about five minutes while the system powers up and pauses for dial-in via the modem. Because it is expecting a dial-in, the system will pause at E07A long enough to allow a remote user to connect to the modem.
4. Depending on whether a password has been set and which password has been entered, either the Main Menu (privileged access password) or General User Menu (general access password) will appear.

B.2 Minimum service processor configuration

The default service processor values are mostly disabled or unconfigured. In order to use the service processor effectively, you need to change the settings to reflect the system behavior you are attempting to control.

The main system behaviors an administrator might want to control include:

- Password protection for the service processor menus
- Operating system surveillance to monitor for hangs
- Call-Out policy to report errors
- Call-In policy to allow or disallow remote support
- Determining power and reboot policies
- Determining firmware levels and learning how to update them

In addition, the system administrator should know how processors can be automatically and manually deconfigured and how to cause the system to boot into the Systems Management Services Menu without having to wait for just the right moment to stop the boot process.

B.2.1 Password protection for service processor menus

Be sure to set the root password on your server. For your own protection, you should also provide password protection for the service processor and Systems Management Services (SMS) menus.

There are two types of passwords for the service processor and SMS. One is a general access password that enables a user to look at Vital Product Data (VPD), service processor and POST error logs, and boot progress indicators. Additionally, the user can continue the current boot and have the progress indicators displayed on the local or remote terminal being used. The second type of password is for privileged users, such as the system administrator. It controls access to the Main Menu of the service processor and SMS utilities.

Passwords can be any combination of up to eight alphanumeric characters. You can enter longer passwords, but the entries are truncated to include only the first eight characters.

The privileged access password can be set from the service processor menus or from SMS utilities. The general access password can only be set from the service processor menus.

For security purposes, the service processor counts the number of unsuccessful attempts to enter a password. After three unsuccessful attempts, the service processor responds in one of two ways depending on whether the attempt was made locally or remotely.

If the attempt was local, the service processor resumes the initial program load (IPL). This is based on the assumption that the server is in an adequately secure location to which only authorized users have access. Users must still successfully enter a login password to access AIX.

If the attempt was via modem, the service processor powers the server down in order to prevent potential security attacks by unauthorized remote users.

Table 16 on page 98 illustrates the impact password settings have on which menu you can access on the service processor. Pay special attention to the so-called *law of unintended consequences*: If you set password protection on

the General User Menu but not the Main Menu, you will have allowed anyone *not* entering a password to access the system's most sensitive settings.

Table 16. Privileged and general access passwords

Privileged Access Password	General Access Password	Resulting Menu
None	None	Service processor Main Menu
None	Set	Users with password see General User Menu; those without go to the Main Menu
Set	None	Users with password see Main Menu, and users without password see General User Menu
Set	Set	Users with privileged access password see Main Menu; users with general access password see General Users Menu

To set passwords from the service processor menus, perform the following steps:

1. From the service processor Main Menu, select **Service Processor Setup Menu**.
2. Select **Change Privileged Access Password**. You should use an eight-character alphanumeric string that cannot be easily guessed.
3. Select **Change General Access Password**. You should use a different eight-character alphanumeric string that cannot be easily guessed.
4. Return to the Main Menu and exit in order to continue the IPL.

B.2.2 Operating system surveillance

The service processor maintains surveillance of the hardware as long as the system is powered on. Using a utility called Repeat-Gard, the firmware automatically marks processors for deconfiguration at the next boot if they experience run-time permanent failures or more than a threshold number of temporary errors. Some errors occur during run-time, but are undetectable by BIST and POST routines. Repeat-Gard will deconfigure the components until they have been replaced. This will prevent getting into a cycle of halting and

rebooting, which would interrupt system availability. If a processor is deconfigured, it remains off-line for subsequent reboots until it is replaced.

In addition to monitoring hardware, the service processor can be configured to keep track of the operating system by means of a heartbeat. By default, this function is disabled.

To enable it, follow these steps:

1. Access the service processor's Main Menu.
2. Select the **Service Processor Setup Menu**.
3. Select **OS Surveillance Setup Menu**.
4. The current setting will be displayed. You may toggle surveillance on or off by pressing the **1** key.
5. Set the surveillance time interval by pressing **2** and entering a value in minutes. The interval is the length of time between heartbeats.
6. Set the surveillance delay in minutes. This is the length of time after the operating system has started before the first heartbeat should be heard.

If the service processor does not hear the heartbeat, it assumes an operating system hang has occurred, and it will shut the system down and reboot. For this reason, the system administrator should carefully determine the time interval of the heartbeat and the delay that might be required to avoid a false hang while applications and resources are being brought up.

B.2.3 Configuring call-out

A basic decision needs to be made on whether to allow the system to call out in the event of a failure. This will require a modem attached to either S1 or S2. Enabling call-out does not mean that the system will automatically support dial-in.

When dialing out, the service processor can either contact IBM, a customer help desk or a digital pager. IBM and the customer help desk must have programs that know how to decode the information that the service processor sends after it connects to a remote modem. IBM Service uses such a program in the United States and many geographies, but you should check locally to see whether this function is available.

If the service processor calls IBM, an electronic problem report will be opened in RETAIN. If the customer's help desk receives the call, a so-called *catcher program* must decode the information. Sample catcher code can be found in /usr/samples/syscatch， but customers will need to build their own

applications to use this function. If a message is sent to a digital pager, only a phone number will be transmitted. By default, the phone number is the one defined as the system, or modem, telephone line. If dial-in is not permitted, another number could be entered into that field. Whoever is carrying the pager would be responsible for knowing which machine the page came from and what steps to take next.

To enable the call-out feature, you need to configure the modem, configure the serial ports, configure the speed of the serial ports, configure the telephone numbers, determine call-out policy, and enter account information.

B.2.3.1 Configuring the modem

Perform the following steps to configure the modem:

1. Have a modem connected to serial port S1 or S2.
2. Access the service processor Main Menu and select the **Call-In/Call-Out Setup Menu**.
3. Select the **Modem Configuration Menu**.
4. Indicate which port will have a modem attached to it and which modem configuration file will be used.
5. After you have linked the correct file to the correct port, save the settings to NVRAM, and configure the modem.
6. Enter `98` to return to the Call-In/Call-Out Setup Menu.

B.2.3.2 Configure the serial ports

Perform the following steps to configure the serial ports:

1. Select the **Serial Port Selection Menu**.
2. Enable Call-Out on the appropriate serial ports.
3. Enter `98` to return to the Call-In/Call-Out Setup Menu.

B.2.3.3 Configure the serial port speed

Perform the following steps to configure the serial port speed:

1. Select the **Serial Port Speed Setup Menu**.
2. Set the speed of the appropriate ports. The default is 9600 baud, which is the slowest recommended speed. Terminal and modem capabilities dictate how much faster you can go. When finished, enter `98` to return to the Call-In/Call-Out Setup Menu.

B.2.3.4 Configure the telephone numbers

Perform the following steps to configure the telephone numbers:

1. Select **Telephone Number Setup Menu**.
2. Assign numbers, appropriate for your geography, for the IBM Service center. If the customer plans to run his own error notification program, configure a number for the help desk. Also, provide entries for Customer Voice Telephone and Customer System Telephone numbers.

> **Note**
>
> If you plan to test call-out, you should initially assign your own telephone number in place of the service center and help desk numbers. That way, you can avoid interrupting ongoing support operations.

3. If you plan to send alerts to a pager, key in the entire string you want to send including the outside access number (if any), the pager service, the recipient's personal identification number, and any intermediate numeric responses needed before the call-back number. The call-back number is the Customer Voice Telephone number. Be sure to insert pauses to allow for delays caused by prompts from the voice response unit.
4. When finished, enter `98` to return to the Call-In/Call-Out Setup Menu.

B.2.3.5 Set the call-out policy

Perform the following steps to set the call-out policy:

1. Select **Call-Out Policy Setup Menu**.
2. Set the policy to either first or all. The service processor will either stop dialing after the first successful connection, or it will dial all three numbers before quitting. You can limit the number of retries.
3. Enter `98` to return to the Call-In/Call-Out Setup Menu.

B.2.3.6 Customer account information

Perform the following steps for customer account information:

1. Select the Customer Account Setup Menu
2. Enter the customer account number in the first field.
3. If applicable, enter the customer's RETAIN login user ID and password.
4. Enter `98` and return to the Call-In/Call-Out Setup Menu.

B.2.3.7 Test the call-out function

Make sure you have changed the IBM Service and customer help desk numbers to point to your own telephone. If you have a digital pager, make sure the whole string has been entered in the Telephone Numbers Menu.

From the Call-In/Call-Out Setup Menu, select **Call-Out Test**. This will generate a pseudo-error that will trigger a call-out to the first or all of the numbers you have entered, depending on the policy you chose to implement.

B.2.4 Configuring call-in

Configuring call-in is simple once you have entered the parameters for the modems and serial ports.

Dial-in must be enabled on the port to which the modem is attached. To test dial-in, shut down the server and perform the following steps:

1. From any telephone, call the server's telephone number. After you hear three rings, hang up. The server powers on.
2. Give the server five minutes to boot up and prepare to receive another call.
3. From an ASCII terminal or terminal emulator, call the server again. The server answers and presents the Service Processor Menus on your terminal.
4. If required, enter your privileged access password. If no password is required, the Main Menu displays.
5. From the Main Menu, select Continue System Boot to view the IPL progress messages. Depending on your server's configuration, the bootup sequence may take several minutes. Once the bootup completes, the logon prompt displays.

You have successfully called into the service processor and brought up the server. Log in and then log out to disconnect from the operating system. Call your server again. The operating system answers and offers the logon prompt.

If these tests are successful, call-in is working correctly.

B.2.5 Power and reboot issues

The S-Series models support unattended start mode and a utility that lets them automatically bring up the operating system after various kinds of failures.

B.2.5.1 Unattended start mode

Unattended start mode means that the service processor automatically restores the system power setting after a temporary power failure. It is intended to be used on servers that require automatic power-on after a power failure.

B.2.5.2 Reboot/Restart Setup Menu options

The service processor will attempt to reboot the system following a run-time failure. You can control some aspects of that behavior. The options may be set by performing the following steps:

1. From the Main Menu, select the **System Power Control Menu**.
2. From the System Power Control Menu, select **Reboot/Restart Setup Menu**.
3. Select **Number of Reboot Attempts** and enter a value.
4. Select **Use OS-Defined Restart Policy** and toggle the value to give the desired setting. If this is set to Yes, which is the default, the operating system will control whether or not to reboot the system following a system crash. If the value is No, the service processor will determine what to do if the system loses control.
5. Select **Enable Supplemental Restart Policy**. The default setting is No. If the operating system has no automatic restart policy, or, if it is disabled, this policy will enable reboots/restarts following a hardware or software failure.
6. Select **Call-Out Before Restart**. By selecting the number of this task, you can toggle between enable and disable. If call-out is enabled, a system fault or surveillance failure will cause the service processor to dial the configured phone numbers using the current call-out policy.
7. When finished, exit from the menus.

B.2.5.3 Saving customized settings

All of the settings you make (except language) from the service processor menus can be saved and used to recover from a fault or to replicate settings to another system that uses the same service processor firmware.

Under the diagnostics service aid Save or Restore Hardware Management Policies, you can save your settings to a file. It is strongly recommended that this service aid be used to protect the usefulness of the service processor and the availability of the server.

B.2.6 Useful service processor utilities

Occasionally, for testing purposes, a system administrator may wish to manually deconfigure one or more processors so that, on the next boot, they will not be initialized. This can be done using a utility accessed through the Main and System Information menus.

B.2.6.1 Configuring/deconfiguring processors

Perform the following steps to configure/deconfigure processors:

1. From the service processor Main Menu, select **System Information Menu**.
2. Select **Processor Configuration/Deconfiguration Menu**.
3. Pick the processor or processors that you want to configure or deconfigure, and toggle them on or off. Hexadecimal numbers following each processor denote its status:
   - 0x00 or 0xFF describe processors configured by the system.
   - 0x81 is for a processor that has been deconfigured manually.
   - 0x41 is deconfigured by the system because the threshold for recoverable run-time errors was exceeded.
   - 0x21 denotes deconfigured due to repeated fatal internal errors.
4. Commit the changes and return to the Main Menu.

The system also uses a utility called Repeat-Gard to deconfigure processors that experience either permanent failures or too many temporary errors. The service processor deconfigures the processor on the next reboot, which, if autorestart is turned on, would follow the system fault. The processor will remain deconfigured until it is replaced.

B.2.6.2 Boot Mode Menu

On the 19990412 version of service processor firmware for S-Series servers, there is a new firmware option that will control service mode boots. Perform the following steps:

1. From the System Power Control Menu, select **Boot Mode Menu**.
2. The first option is to boot to the SMS Menu. You can toggle this value on and off. If this value has been toggled on, you do not need to manually wait for the right moment to press the terminal enter keys in order to bring up the SMS menu.
3. The next option is Service Mode Boot from Saved List, which boots from the device specified in the service mode boot list saved in NVRAM. If AIX diagnostics has been installed, AIX boots in single-user mode to the diagnostics menu. Using this option to boot the system is the preferred method for running online diagnostics. In this mode, diagnostics has access to the system error log.
4. The next option is Service Mode Boot from Default List, which boots from the devices as defined in system firmware. This is the preferred way of

booting from a CD into standalone diagnostics. In this mode, diagnostics does not have access to disk or the system error log, but can test devices that would be locked if AIX were running.

5. The final option is Boot to Open Firmware Prompt. This option should only be used when you are directed to do so by support personnel or the vendor of a third-party device that uses Open Firmware for installation.

B.3 Determining firmware levels

If AIX is running, the system and service processor firmware levels can be determined using the `lscfg` command. If the server is about to be initialized, the firmware levels can be viewed using the SMS utilities for the system and the service processor menus for the service processor.

Use one of the following procedures to determine your firmware level:

At the AIX command prompt on your system, enter the following command:

```
lscfg -vp | grep -p alterable
```

This command will produce a system configuration report similar to the following:

```
System Firmware:
ROM Level.(alterable).......19990629 (B) 19990621 (A) <= System FW Levels
  Version.....................RS6K
System Info Specific.(YL)...P2
SP_CARD_:
  Part Number.................PART_NUM
  EC Level....................EC_LEVEL
  FRU Number..................FRU_NUM_
  Manufacture ID..............IBM
  Serial Number...............SERIAL_#
  Version.....................0000RS6K
ROM Level (alterable).......19990630 (B) 19990620 (A) <= SvP FW levels
System Info Specific.(YL)...P2
```

The lines that start *ROM Level (alterable)* list the level numbers of the installed system and service processor (SvP) firmware. The levels match the date on which the update was released, in the format YYYYMMDD.

In these examples, the system was booted from type B because it is listed first. The system firmware is, therefore, 19990629. A slightly older copy of the system firmware is in type A, 19990621. You should consider updating your system firmware when new releases become available. When an update is

performed, the firmware is written to type B only and can be promoted to type A at a later time. The system firmware level can also be seen from System Management Services menus.

In the above example, the service processor firmware in use is 19990630. Another way of determining this level is to access the service processor menus. The firmware level is contained in the heading of the main menu. The following is an example of what the heading looks like:

```
Service Processor Firmware
    Version: 19990630
Copyright 1998, IBM Corporation
```

The numbers in the second line show the booted firmware level, 19990630. If this level is less than the update level available for your server, you should consider installing the update.

To determine the level of your system firmware via Systems Management Services, perform the following steps:

1. Power the system on after it has been shut down, or reboot the system.

2. Watch for the character-based RS/6000 logo screen and the POST indicators to appear on your terminal.

3. When the word Keyboard appears, immediately press the **1** key. The number 1 key must be pressed before the word Speaker appears.

When the tests have completed and any required passwords have been entered, the System Management Services Utilities menu appears. The System Firmware level is displayed in the top left-hand corner of the display.

When you have read the current firmware level, exit the System Management Services menu as directed on the screen.

If you find the firmware level is less than the update level available for your server, you should consider installing the update.

B.4 Updating firmware

Firmware update packages are available from the RS/6000 Support page on the Internet under RS/6000 Microcode Updates. The URL is:

`http://www.rs6000.ibm.com/support/`

The system firmware and service processor firmware are combined into a single download package. There is one package for the Model S70 and Model S70 Advanced and another package for the Model S80.

Prior to downloading the firmware, you are asked to read and accept the terms of the Machine Code License Agreement. Once you accept the terms, you are assigned a password. Write down this password because it is required later to unpack the files you download.

Find the most recent update package for your server. Print the description file, and download one of the format choices, depending on the workstation being used for downloading. The description file provides detailed instructions on downloading and updating the firmware.

Be certain that you do not attempt to load firmware that is inappropriate for your system. Make sure that you follow instructions to the letter including directions on which module to install first - system or service processor. Be aware that installation instructions may vary depending on the level of firmware. Always read the README file before attempting to update firmware.

Appendix C. Special notices

This publication is intended to help IBM and Business Partner sales and technical support staff understand the architecture, features, and benefits of the RS/6000 S-Series Enterprise Servers. The information in this publication is not intended as the specification of any programming interfaces that are provided by the AIX Operating System, Program Number 5765-C34. See the PUBLICATIONS section of the IBM Programming Announcement for the AIX Operating System, Program Number 5765-C34, 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 the IBM product, program, or service may be used. Any functionally-equivalent program that does not infringe any IBM intellectual property rights may be used instead of the IBM product, program, or service.

Information in this book was developed in conjunction with the 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:

AIX
AS/400
ESCON
IBM
IBMLink
LANStreamer
Predictive Failure Analysis
RETAIN
RS/6000
S/390
SecureWay
Service Director
SP
Streamer
400

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/or 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 D. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 IBM Redbooks publications

For information on ordering these publications see “How to get IBM Redbooks” on page 115.

- *AIX Version 4.3 Differences Guide*, SG24-2014

D.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

D.3 Other resources

These publications are also relevant as further information sources:

- *PCI Adapter Placement Reference*, SA38-0538
- *RS/6000 Enterprise Server S70 Advanced Installation and Service Guide*, SA38-0548
- *RS/6000 Enterprise Server S80 Installation and Service Guide*, SA38-0558
- *Site and Hardware Planning Information*, SA38-0508

- *RS/6000 SP Planning Volume 1, Hardware and Physical Environment*, GA22-7280

D.4 Referenced Web sites

The following Web sites are also relevant as further information sources:

- http://www.rs6000.ibm.com/resource/hardware_docs/
- http://www.rs6000.ibm.com/library/
- http://www.rs6000.ibm.com/support/micro
- http://www.tpc.org
- http://www.spec.org

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.ibmlink.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.ibmlink.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.ibmlink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at http://w3.ibm.com/ for redbook, residency, and workshop announcements.

IBM Redbooks fax order form

Please send me the following:

Title	Order Number	Quantity

First name	Last name	
Company		
Address		
City	Postal code	Country
Telephone number	Telefax number	VAT number

☐ Invoice to customer number _____

☐ Credit card number _____

Credit card expiration date	Card issued to	Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Glossary

ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BI	Business Intelligence
BIST	Built-In Self Test
CEC	Central Electronic Complex
CPU	Central Processing Unit
CRC	Cyclic Redundancy Checking
CWS	Control Workstation
ECC	Error Checking and Correction
ERP	Enterprise Resource Planning
ESCON	Enterprise Systems Connection
FC-AL	Fiber Channel-Arbitrated Loop
FDDI	Fiber Distributed Data Interface
HACMP	High Availability Cluster Multi Processing
HTTP	Hypertext Transfer Protocol
IA	Intel Architecture
IA-32	32-bit Intel Architecture
IA-64	64-bit Intel Architecture
IBM	International Business Machines Corporation
I/O	Input/Output
IP	Internet Protocol
ITSO	International Technical Support Organization
JFS	Journal File System
L1	Level 1
L2	Level 2
LCD	Liquid Crystal Display
LPAR	Logical Partitioning
LVD	Low Voltage Differential
LVM	Logical Volume Manager
MES	Miscellaneous Equipment Specification
NIM	Network Install Manager
NIS	Network Information System
NVRAM	Non-Volatile Random Access Memory
OLTP	On-line Transaction Processing
PCI	Peripheral Component Interconnect
POST	Power-On Self Test
PSSP	Parallel System Support Program
RAID	Redundant Array of Independent Disks
RAS	Reliability Availability Serviceability
RIO	Remote Input/Output
SCSI	Small Computer System Interface
SDRAM	Synchronous Dynamic Random Access Memory
SE	Single Ended
SMIT	System Management Interface Tool

SMP	Symmetric Multiprocessor
SPEC	System Performance Evaluation Corporation
SSA	Serial Storage Architecture
TCP	Transmission Control Protocol
TPC	Transaction Processing Performance Council
VHDCI	Very High Density Cable Interconnect
VMM	Virtual Memory Manager
VPD	Vital Product Data
WLM	Workload Manager

Index

Symbols

Numerics

A

B

C

J

K

L

M

N

O

P

Q

R

S

IBM Redbooks evaluation

RS/6000 S-Series Enterprise Servers Handbook
SG24-5113-00

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!)

_____

_____

_____

_____

SG24-5113-00
Printed in the U.S.A.

RS/6000 S-Series Enterprise Servers Handbook

SG24-5113-00