IBM @server pSeries 660 Model 6M1

Technical Overview and Introduction

September 4, 2001

Stephen Lutz

Shyam Manohar

Scott Vetter

International Technical Support Organization

Austin, TX

IBM

IBM @server pSeries™ 660 Model 6M1

The IBM @server pSeries™ 660 Model 6M1 (referred to hereafter as the Model 6M1) was announced on September 4, 2001. The Model 6M1 shares its heritage with the IBM @server pSeries 680, which achieved leadership performance in the high-end SMP UNIX server marketplace with a TPC-C benchmark result of 220,807 tpmC at $34.18 per tpmC[1]. This provides the best price/performance of any non-clustered server to exceed 70,000 tpmC. It also recorded the highest SpecWeb99 rating of 9106 simultaneous connections[2].

The outstanding performance of the IBM @server pSeries 680 was achieved through a combination of balanced systems design and unique IBM Silicon-on-Insulator technology. The Model 6M1 enhances the mid-range server lineup with many of the same design elements.

Overview

The Model 6M1 introduces additional features beneficial for a Model M80 customer wishing to upgrade, such as Capacity Upgrade on Demand, improved throughput, and faster processors. Existing Model M80 customers may select memory and processor features of the Model 6M1, effectively upgrading their systems to Model 6M1 performance levels.

The Model 6M1 provides the maximum performance scalability and I/O expandability among the RS/6000 and pSeries midrange servers by supporting the greatest number of processors and PCI slots.

This paper discusses, in detail, the processor, memory, I/O, expandability, reliability, and other technical aspects related to the Model 6M1.

pSeries 660 Model 6M1 Introduction

The Model 6M1 is a member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM that support a range of 32- and 64-bit applications simultaneously.

The Model 6M1 provides scalability using the 64-bit RS64 IV processor packaged in 2-, 4-, 6-, and 8-way 750 MHz SMP configurations, from 2 GB to 64 GB of real memory and incorporates an I/O subsystem supporting 32-bit and 64-bit standard PCI adapters. A 2-way or 4-way 500 MHz RS64 III SMP configuration is also available for initial orders. 2-way 500 MHz systems may add a second 2-way 500 MHz processor card, or replace the existing processor cards with any of the 750 MHz processor features (see “Clock, Power, and Backplane” on page 8 for additional information) to meet your computing needs.

Central Electronics Complex and I/O Drawer

The Model 6M1 package consists of a central electronics complex (CEC) and I/O drawers. The CEC drawer incorporates the system processors, memory, and supporting system logic.

[1] per: www.ideasinternational.com as of 4/13/2001
[2] per: www.spec.org as of April 2001

The CEC is 8 U (EIA units)[3] high. The dimensions are 445 mm W x 826 mm D x 356 mm H (17.5” W x 32.5” D x 14.0” H). The weight of the CEC is from 59.7 kg (132 lbs) to over 74.6 kg (169 lbs), depending on the configuration.

The I/O drawers are connected with different cables to the CEC (see “Cabling” on page 16). Each I/O drawer provides 14 hot-plug PCI slots. The primary I/O drawer also provides connectors for external devices and the service processor. A maximum of four I/O drawers, with up to 56 hot-plug PCI slots, are supported in a fully featured Model 6M1 system.

The I/O drawers are 5 U in height each. The dimensions are 445 mm W x 820 mm D x 218 mm H (17.5” W x 32.3” D x 8.6” H). The weight of one I/O drawer is from 41 kg (90 lbs) to 52 kg (115 lbs) depending on the configuration.

The Model 6M1 is designed to be installed and maintained by trained service representatives. All adapters and devices that are part of the original order are installed and configured before shipment.

The drawers can be mounted in an existing 7014 Model S00, Model R00, or in the available 7014 Model T00 and 7014 Model T42 racks. The recommended rack is the 7014 Model T00 (see “Enterprise Racks” on page 3).

Figure 1 shows the front view of a Model 6M1 with two I/O drawers and an external storage enclosure in a T00 rack.

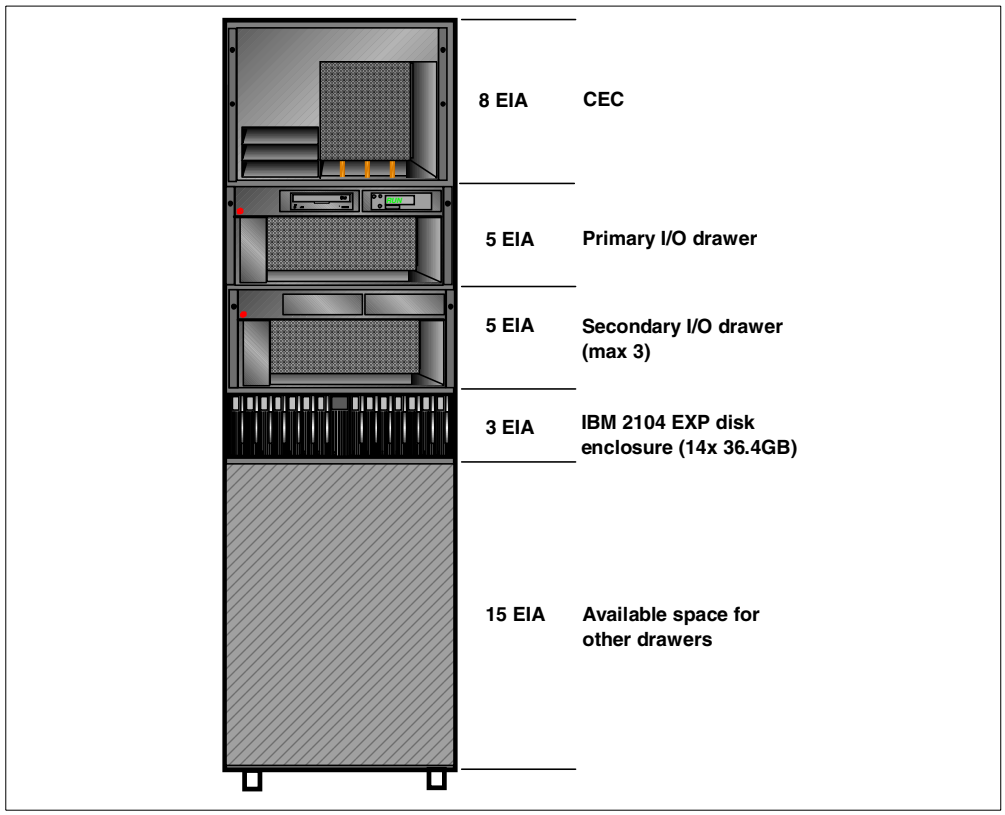

Figure 1. Model 6M1 with Two I/O Drawers and Additional Storage in T00 Rack

[3] One EIA unit (Electronic Industries Association Unit) is 44.45 mm or 1.75”

Internal Storage

The system comes preconfigured with a CD-ROM and a diskette drive mounted in the primary I/O drawer, leaving one media bay free for customer expansion, such as a tape device (disk drives are not supported in the media bay). A DVD-RAM is an available feature that can be used as a writable backup device as well as for reading installation media. Any devices in the media bays of the primary I/O drawer are connected to the internal F/W SCSI controller (no additional cable is required). External disk storage is required for the Model 6M1 as data storage; however, optional boot bays are available in the primary I/O drawer. For further information concerning storage, refer to “I/O Subsystem Architecture” on page 10.

Note

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

Power and Cooling Options

In a minimum configuration, both the CEC drawer and the primary I/O drawer are equipped with redundant power supplies. These power supplies can be either AC (# 9172) or -48V DC (# 9175), and both are designed to support up to 8-way SMP systems with 64 GB of memory. For secondary I/O drawers, additional redundant power supplies (# 6283) are required and must be part of the configuration. Concerning DC power, a redundant power supply is also standard in the CEC and in all I/O drawers. Therefore, redundant power supplies do not need to be added separately. All power supplies are hot-plug and allow concurrent repair.

In each system, a processor power regulator must be ordered. Power regulators are used to transform and filter the raw voltage feeds into voltages usable by the chips and components. For a 2- or 4- way 500 MHz RS64 III system, a 76A power regulator (# 6196) is required. For 2-, 4-, 6-, or 8-way systems with 750 MHz RS64 IV processors, a different regulator (# 6164) is required.

The Model 6M1 CEC and I/O drawer cooling subsystems have N+1 hot-plug fans, which are monitored for performance by SPCN (see “System Power Control Network (SPCN)” on page 20). In the event of a single fan failure, the other fans increase their speed to provide sufficient cooling. The hot-plug fans can be replaced concurrently.

Enterprise Racks

This section provides an overview of the Enterprise racks that are available from IBM.

Enterprise Rack Model T00 and T42 are 19-inch wide racks for general use with pSeries and RS/6000 rack drawer systems.

IBM RS/6000 7014 Model T00 Enterprise Rack

The 1.8 meter (71 inch) Model T00 is compatible with past and present RS/6000 and pSeries racks, and is designed for use in all situations that have previously used the older rack Models R00 and S00. The improved features in the T00 rack are as follows:

- 36 EIA units (36 U) of usable space.
- Optional removable side panels.

- Optional front door.
- Optional side-to-side mounting hardware for joining multiple racks.
- Increased power distribution and weight capacity.
- Standard black or optional white color.
- Optional reinforced (ruggedized) rack feature (# 6080) provides added earthquake protection with modular rear brace, concrete floor bolt down hardware, and bolt-in steel front filler panels.
- Model T00 supports both AC and DC configurations.
- Weight:
  - T00 base empty rack: 244 kg (535 pounds)
  - T00 full rack: 816 kg (1,795 pounds)

IBM RS/6000 7014 Model T42 Enterprise Rack

The 2.0 meter (79.3 inch) Model T42 is the rack that will address the special requirements of customers who want a very tall enclosure to house the maximum amount of equipment in the least amount of floor space. The Model T42 rack has all of the enhancements found in the Model T00 with the following exceptions:

- 42 EIA units (42 U) of usable space
- Model T42 supports AC only
- Weight:
  - T42 base empty rack: 261 kg (575 pounds)
  - T42 full rack: 930 kg (2,045 pounds)

Rack-Mounting Rules for Model 6M1

There are some rules that should be considered when mounting drawers into a rack:

- The Model 6M1 I/O drawers are designed to be placed at any location in the rack.
- The Model 6M1 CEC can be mounted at the top of the rack, since it has rear access for service.
- Any space in the rack can be used for storage, such as 7133-D40, if desired.

A Model 6M1 with one I/O drawer is 13 U in height, so a maximum of two Model 6M1s fit in an S00 or T00 rack, or you can put three Model 6M1s in a T42 rack.

IBM 7316-TF1 Rack Console

New for rack-mounted systems is the ability to install a system console closely to the hardware mounted in a system rack. This monitor combines a bright 15-inch viewable image and 1024 x 768 addressability with a space-saving package design. Using thin film transistor (TFT) LCD technology, the T54A features a 15 inch (304.1 mm x 228.1 mm) viewable area with flicker-free display of the primary XGA-mode (1024 x 768) and full-screen support for other common Video Electronics Standards Association (VESA) and industry modes. The 7316-TF1 Rack Console has the following attributes:

- 3 EIA units (3 U)
- IBM T54A Flat Panel Monitor

- Flat panel monitor rack-mounted kit
- Rack keyboard tray
- Optional IBM Space Saver 2 Keyboard that mounts in the Rack Keyboard Tray (track point mouse is integral to keyboard)

Note

The 7316-TF1 requires a GXT135P Graphics Adapter (#2848) or equivalent graphics adapter.

VGA Switch

The VGA Switch for the IBM 7316-TF1 Rack Console allows users to control multiple servers from a single console. This dual-user switch allows attachment of one or two consoles, one of which must be an IBM 7316-TF1. Either console can control any one of the eight servers, except that both consoles cannot talk to the same server at the same time. An easy to use graphical interface, supported in six languages (English, French, Spanish, German, Italian, Brazilian Portuguese), allows fast switching between systems. Using multiple switches in a two-level cascade arrangement, as many as 64 systems can be controlled from a single point.

The VGA Switch is only 1 EIA unit (1U) high and can be mounted in the same tray as the 7316-TF1 Rack Console, thus conserving valuable rack space. It supports a maximum video resolution of 1600 x 1280, which facilitates the use of graphics-intensive applications and large monitors.

To help minimize cable clutter, multi-connector cables in lengths of 7, 12 and 20 feet are available. These cables can be used to connect the graphics adapter (required in each attached system), keyboard port and mouse port of attached servers to the switch or to connect between switches in a tiered configuration.

System Architecture and Technical Overview

This section introduces the technical aspects of the CEC and I/O drawer.

CEC Architecture

Figure 2 shows the system schematic of the Model 6M1 with RS64 IV processors within the CEC. This section discusses the different physical components and the SMP processor configurations in the system schematic.

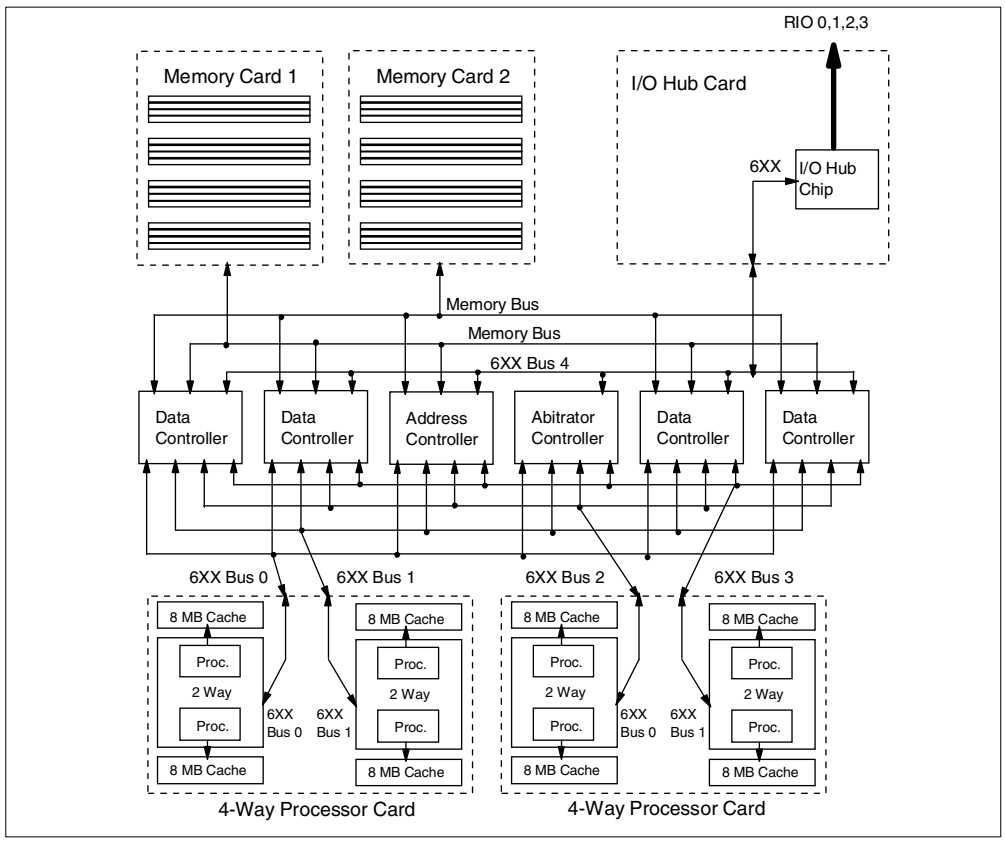

Figure 2. RS/6000 Model 6M1 System Schematic within the CEC

RS64 III RISC Processor

An RS64 III processor card used in entry level Model 6M1 (# 5200) configurations provides a 2- or 4-way SMP with the following attributes:

- 500 MHz operating frequency
- CMOS 7S manufacturing process
- 128 KB on-chip L1 instruction cache with parity and prefetch
- 128 KB on-chip L1 data cache with ECC
- On-chip L2 cache directory
- 4 MB per processor of off-chip L2 cache using ECC double data rate (DDR) SRAM
- PowerPC 6xx bus architecture, 16-byte wide bus interface

RS64 IV RISC Processor

The RS64 IV processor card used in the Model 6M1 (# 5210, # 5213, # 8306, # 8307) provides 2-, 4-, 6-, or 8-way configurations with the following attributes:

- 750 MHz operating frequency
- CMOS 8S2 manufacturing process
- 128 KB on-chip L1 instruction cache, 2-way associative, with parity and retry for data integrity
- 128 KB on-chip L1 data cache, 2-way associative, with ECC
- On-chip L2 cache directory
- 8 MB of off-chip L2 cache using ECC double data rate (DDR) SRAM per processor
- PowerPC 6xx bus architecture, 16-byte wide bus interface

The RS64 IV processor in the Model 6M1 has an operating frequency of 750 MHz. The frequency is accomplished by leveraging the IBM copper (CMOS 8S2) and silicon-on-insulator (SOI) technology.

The copper technology and an improved manufacturing process allow the chip to operate at 1.8V. The lower operating voltage, coupled with the smaller circuit dimensions, result in reduced wattage in the RS64 IV and allows additional function to be placed on the chip.

SOI technology reduces signal loss due to capacitance that normally occurs along the silicon/conductor boundaries on the chip using a thin insulating boundary. This improves signal-to-noise ratios within the chip and provides additional freedom in the layout of the circuit design.

The size of the level one (L1) instruction and data caches is 128 KB each. Innovative custom circuit design techniques were used to maintain the one cycle load-to-use latency for the L1 data cache. The level two (L2) cache directory was integrated into the RS64 IV chip, reducing off-chip accesses that impact performance.

IBM uses double data rate (DDR) SRAM technology for the L2 cache in the RS64 IV processor. DDR technology provides two transfers of data on the 16-byte wide L2 data bus every SRAM clock cycle. The DDR SRAM technology also reduced L2 access latency as measured in nanoseconds. This technology allows L2 accesses to perform at processor clock speed.

Hardware Multi-Threading (HMT)

The RS64 IV supports hardware multi-threading (HMT) wherein two logical processors share one physical hardware processor. Hardware multi-threading provides a mechanism of improving overall system throughput in specific cases by overlapping memory access with other computation. AIX® Version 4.3.3 and AIX 5L™ Version 5.1 support HMT. It is enabled by the `bosdebug -h on` command and activated on subsequent reboot. When invoked, AIX presents the system as having twice as many processors as are physically installed. Each has typically a little more than one half the performance of a non-HMT system. For further information, see /usr/lpp/bos/README.HMT as shipped with the AIX operating system. This feature cannot be used if Dynamic Processor Deallocation (“Dynamic Processor Deallocation” on page 19) is active. Application

benchmarking of production workloads on test machines is recommended to determine if the workload will benefit from HMT.

Processor Boards

The CEC drawer backplane in the Model 6M1 consists of the memory controller chip set and includes slots for processor boards, power regulator, clock card, memory cards, and an I/O card. The number of slots available for the processor boards in the backplane is two. There are three types of processor boards that can be used independently, or in combination, to build 2-, 4-, 6-, and 8-way SMP configurations in the Model 6M1, which are explained as follows:

2- and 4-Way SMP

The Model 6M1 system can be configured as a 2-way SMP configuration using a single 750 MHz RS64 IV (# 5210) or 500 MHz RS64 III (# 5200) board. The Model 6M1 system can be configured to a 4-way SMP by adding one more dual-processor board on the other processor board slot, available on the CEC backplane.

4-Way SMP

A single four-processor board can be used to configure 4-way SMP configuration in the Model 6M1 (# 5213). Adding the Capacity Upgrade on Demand feature (# 8306 for zero additional processors active, or # 8307 for two additional processors active) allows a 4-way board to extend the number of processors from four up to eight. This feature requires a special licensing agreement to be signed with IBM and an existing 4-way board (# 5213) to be in place. This board is available only as a 750 MHz RS64 IV board.

6-Way SMP

A 6-way SMP can be configured by using one dual processor board (# 5210) and one four-processor board (# 5213). Both must be 750 MHz RS64 IV boards.

8-Way SMP

As discussed for a 4-way SMP, a single four-processor board can be used to configure 4-way SMP configuration in the Model 6M1. By using an additional four-processor board (# 5213), the Model 6M1 can be configured to an 8-way SMP using 750 MHz RS64 IV processors.

In the Model 6M1, each pair of RS64 III or RS64 IV processors share a single processor bus interface to the memory controller. While a dual processor board requires a single processor bus interface, the four-processor board requires two. The total processor bus interfaces required for 6- and 8-way SMP configuration are three and four, respectively.

Clock, Power, and Backplane

The adaptation of the RS64 IV processor in the CEC of the Model 6M1 required a series of engineering changes to the Model M80 design that includes a new clock card, power regulator card, and CEC backplane. The new clock card provides a 1:10 clock ratio, with the CPU clock providing a 75 MHz bus speed to memory. It also provides a 1:5 clock ratio with the CPU clock, providing a 150 MHz bus speed on the PowerPC 6xx processor bus.

The power regulator handles the requirements of the RS64 IV processor cards. An improved backplane, though functionally equivalent to a Model M80

backplane, increased the signal quality required to operate at the higher internal speeds.

Model 6M1 systems ordered with 500 MHz RS64 III processors include the improved backplane.

Memory Controller

The memory controller chip set in the Model 6M1 provides five PowerPC 6xx bus interfaces and two high-speed memory interfaces.

The RS64 IV processors use the PowerPC 6xx bus as the processor interface bus. In the Model 6M1, the processors use four out of the five 6xx bus interfaces provided by the memory controller. The I/O hub uses the remaining bus. The PowerPC 6xx bus is a 16-byte wide bus and operates at a clock rate of 150 MHz.

Memory Subsystem

The memory controller chip set in the Model 6M1 provides two high-speed memory bus interfaces and provides the functions of ECC as well as memory scrubbing. Memory scrubbing provides a built-in hardware function that is designed to perform continuous background reads of data from memory, checking for correctable errors. The two memory interfaces are connected to two memory riser cards. In the Model 6M1 with 750 MHz processors, each memory bus that provides the interface between the memory controller and a riser card is 64 bytes wide and operates at a clock rate of 75 MHz.

Each riser memory card provides 32 DIMM slots. The Model 6M1 uses 200-pin 10ns SDRAM DIMMs. The memory is populated in octals on the DIMM slots with DIMMs of equal sizes. DIMM size used in one octal can, however, coexist with a different DIMM size used in another octal. The minimum configuration requires one riser memory card. A riser memory card should have a minimum of one octal of DIMMs populated.

The overall memory bandwidth in the Model 6M1 can be exploited by using both of the riser memory cards and by distributing the memory on the cards. Each riser memory card provides memory expandability of up to 32 GB by populating all the slots using 1024 MB DIMMs. With both the riser memory cards fully populated, the Model 6M1 provides memory expandability up to 64 GB.

Bus Bandwidth

The following are the theoretical maximum bandwidths, as applicable for an 8-way 750 MHz SMP configuration:

- Bandwidth of the bus between each memory riser card to memory controller: 4.8 GB/s
- Bandwidth of the PowerPC 6xx bus used to interface each pair of processors: 2.4 GB/s
- Bandwidth of the PowerPC 6xx bus used to interface the I/O hub: 2.4 GB/s
- Aggregate memory bandwidth: 9.6 GB/s
- Aggregate processor bandwidth: 9.6 GB/s
- Four drawer I/O bandwidth: 4 GB/s (4 x 500 MB/s bidirectional)

I/O Hub Card

One of the five PowerPC 6xx buses provided by the memory controller is interfaced to the I/O hub chip, residing on a separate I/O hub card that is plugged into the CEC backplane. The I/O hub chip provides four Remote I/O (RIO) ports, each consisting of two unidirectional, 1-byte wide links operating at a 500 MHz clock rate. Each pair of RIO ports is used to connect a maximum of two I/O drawers in a loop fashion. The Model 6M1 supports up to four I/O drawers using the two pairs of RIO ports. The connectors for the four RIO ports are provided on the I/O hub card.

I/O Subsystem Architecture

The I/O subsystem consists of a primary I/O drawer and, if more PCI slots are necessary, up to three additional secondary I/O drawers, providing up to 56 hot-plug PCI slots. The I/O drawers in the Model 6M1 are the same as the I/O drawers in the Model 6H1.

The primary I/O drawer contains the following features:

- Fourteen available hot-plug PCI slots per drawer (up to 56 with four I/O drawers)
- Service processor (only primary I/O drawer)
- One integrated SCSI-2 F/W port for the internal drawer components
- One Ultra2 SCSI port for external attachment use (mini 68-pin VHDCI[4] connector)

  The industry-standard VHDCI 68-pin connector on the backside of the I/O drawer allows attachment of various LVD and SE external subsystems. A 0.3 meter converter cable, VHDCI to P, mini 68-pin to 68-pin, (# 2118) can be used with older external SE subsystems to allow a connection to the VHDCI connector.
- CD-ROM drive or optional DVD-RAM
- Diskette drive (mounted behind a removable cover)
- One additional media bay (on the left side of the drawer)
- 10/100 Mb/s Ethernet port (RJ-45 connector)
- Keyboard and mouse port
- Four serial ports (max. 230 Kb/s, 9-pin D-shell)

  Systems include two 9-pin to 25-pin converter cables (equivalent to # 3925)
- One parallel port (bidirectional)

The Debug port is for diagnostics and is normally covered with a metal plate. The Debug port uses the same connector as the parallel port. To avoid confusion, this port should always remain covered.

[4] Very High Density Cable Interconnect (VHDCI)

Figure 3 shows the rear view of the primary I/O drawer and the port locations.

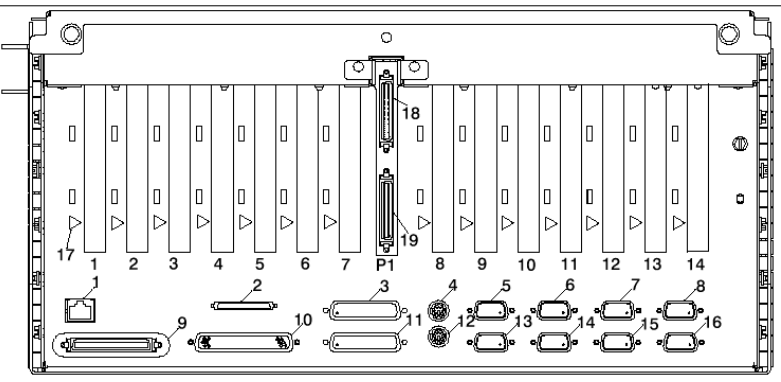

1	Ethernet connector	2	SCSI connector
3	Debug	4	Mouse connector
5	Serial port 1	6	Serial port 2
7	J11 (SPCN 1 connector)	8	J14 (SPCN 3 connector)
9	V/S COMM	10	JTAG connector
11	Parallel port	12	Keyboard connector
13	Serial port 3	14	Serial port 4
15	J15 (SPCN 2 connector)	16	J16 (SPCN 4 connector)
17	Hot-Plug LED (one for each adapter)	18	RIO 1 connector
19	RIO 0 connector		

Figure 3. I/O Drawer Rear View

Operator Panel

The primary I/O drawer incorporates the primary operator panel and indicators for the system. The panel consists of the following features (see Figure 4 on page 12):

- Power On/Off button
- Power-on LED (green)
- System Attention LED (yellow)

  This LED indicates to a user that there is an attention condition on the system.
- Operator panel display

  The display has two lines of sixteen characters each. The display shows reference codes from the service processor, the SPCN, and the operating system. These codes can be either informational codes or error codes. Informational codes will have the System Attention LED off; error codes will have the System Attention LED on.
- Service processor reset button

  The service processor reset button is a single execution button used to reset the service processor and bring the system back into standby mode. Access to this button is restricted through a pinhole in the operator panel cover. This button is for service use only.
- Speaker (beeper)

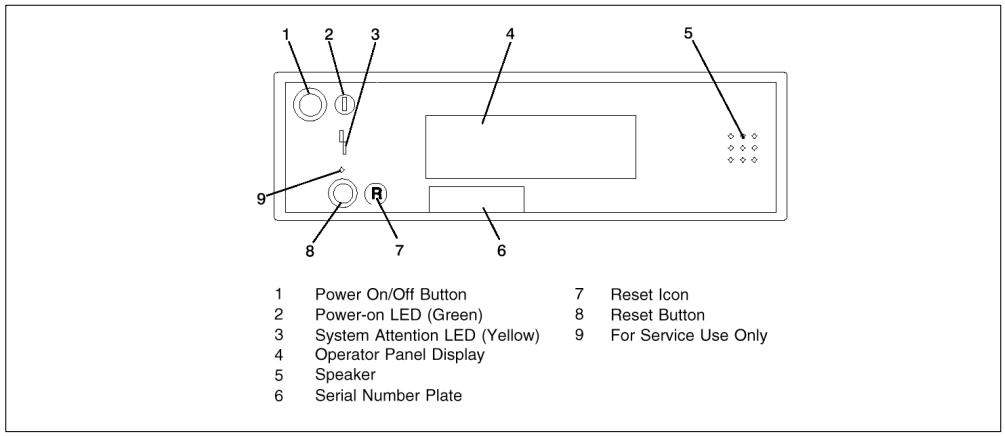

Figure 4. Operator Panel

I/O Drawer Electronics

Figure 5 shows the major electronic components of the I/O drawer. These components are discussed in the following section.

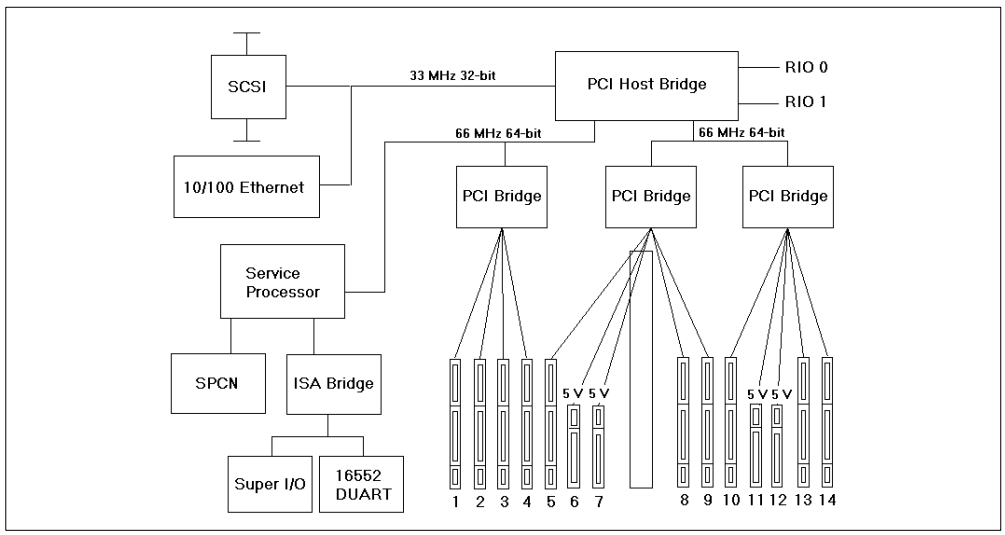

Figure 5. Slot Layout of the I/O Drawer

The drawers are manufactured in a *sandwich architecture*, meaning there are two main circuit boards, one on the bottom of the I/O drawer with the electronics for the integrated ports and service processor, and the board for hot-plug adapter cards on top of the other. This architecture prevents damage to the components in the I/O drawer when plugging hot-plug PCI cards into the system while it is running. Other features make working with your I/O drawer easier, such as quick releases, for easy power supply swaps, and cable routers, to help organize cables on the back of the drawer.

The I/O host bridge in the I/O drawer allows the two RIO ports to connect to the CEC, and three primary PCI buses. One PCI bus operates at 33 MHz, and the two others operate at 66 MHz.

- The first bus is used for the on-board dual SCSI adapter (F/W SCSI internal, Ultra2 SCSI external) and the on-board 10/100 Mbps Ethernet adapter.
- The second bus connects the service processor, SPCN, a PCI-to-ISA bridge chip and a PCI-to-PCI bridge chip. Two I/O chips, a National Super I/O chip and a 16552 DUART[5], are connected to the ISA bus. The Super I/O chip provides the floppy drive controller, two of the four serial ports, keyboard and mouse ports, and the parallel printer interface. The 16552 DUART provides serial ports three and four. The PCI-to-PCI bridge chip provides four 64-bit hot-plug PCI slots.
- The third bus is connected to another two PCI-to-PCI bridges, which provide another six 64-bit and four 32-bit hot-plug PCI slots.

Each slot represents a separate PCI bus, which simplifies the hot-plug function.

PCI Slots

All PCI slots are PCI 2.2 compliant and are hot-plug enabled, which allows most PCI adapters to be removed, added, or replaced without powering down the system. This function enhances system availability and serviceability.

The ten 64-bit slots operate at 3.3V signaling at 66 MHz, which is in contrast to the four 32-bit slots that operate at 5V signaling at 33 MHz (see Figure 5). When adding adapters to the system, it is important which signaling the adapter uses: 3.3V, 5V, or universal, which means the adapter works at both voltages. That is the reason why a PCI 3-Channel Ultra2 SCSI RAID Adapter (# 2494) can be placed only in slots 6, 7, 11, or 12. Refer to the *PCI Adapter Placement Reference Guide*, SA38-0538 for further information.

32-bit versus 64-bit PCI Slots

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

32-bit adapters can typically function in 64-bit slots; however, 32-bit adapters still operate in 32-bit mode and offer no performance advantages in a 64-bit slot. Likewise, most 64-bit adapters can operate in 32-bit PCI slots, but will operate in 32-bit mode at a reduced performance potential.

Hot-Plug PCI Adapters

The function of hot-plug PCI adapters is to provide concurrent adding or removal of PCI adapters when the system in running.

In the I/O drawer, the installed adapters are protected by plastic separators, designed to prevent grounding and damage when adding or removing adapters. The hot-plug LEDs outside the I/O drawer (see Figure 3) indicate if an adapter can be plugged in or removed from the system. These LEDs are also visible inside the I/O drawer. Inside, the light from the LED is routed to the top of the plastic separators using light pipes, which makes it very easy to locate the right slot. The hot-plug PCI adapters are secured with retainer clips on top of the slots, therefore, you do not need a screwdriver to add or remove a card and there is no screw that can be dropped inside the drawer.

[5] Dual Universal Asynchronous Receiver-Transmitter (DUART)

The function of hot-plug is not only provided by the PCI slot, but also by the function of the adapter. Most adapters are hot-plug, but some are not. Be aware that some adapters must not be removed when the system is running, such as the adapter with the operating system disks connected to it or the adapter that provides the system console. Refer to the *PCI Adapter Placement Reference Guide*, SA38-0538 for further information.

To manage hot-plug PCI adapters, it is important to turn off slot power before adding, removing, or replacing the adapter, which is done by the operating system. The following methods can be used to manage hot-plug PCI slots in AIX:

- Command line:
  - `lsslot` - List slots and their characteristics
  - `drslot` - Dynamically reconfigures slots
- SMIT
- Web-based System Manager

When working with the commands and tools mentioned above, the hot-plug LEDs (see Figure 3) change their state. Table 1 shows the possible states of the hot-plug LEDs.

Table 1. Hot-Plug LED Indications

LED Indication	PCI Slot Status	Definition
Off	Off	Slot power is off. It is safe to remove or replace adapters.
On (not flashing)	On	Slot power is on. Do not remove or replace adapters.
Flashing slowly (one flash per second)	Identify	Indicates the slot has been identified by the software; do not remove or replace adapters at this time.
Flashing fast (six to eight flashes per second)	Action	Indicates the slot is ready for adding, removing, or replacing of adapters.

To add a hot-plug PCI adapter, use the `drslot` command to set the slot first into the Identify state (LED flashes slowly) to verify the right slot was selected. After pressing Enter, the LED changes its state to the Action state (LED flashes fast). Then, add the adapter to the system. When finished, press Enter again to turn on the slot power. The hot-plug LED will change its state to On. Now, the adapter is integrated into the system and can be configured using AIX `cfgmgr` (configuration manager).

Before adapters can be removed, they must be deconfigured in AIX. The adapter must be in a defined state or removed from the ODM.

Figure 6 on page 15 shows an example of adding a hot-plug PCI adapter to a running system.

```
# lsslot -c pci
# Slot       Description                       Device(s)
U0.1-P1-I1   PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I2   PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I3   PCI 64 bit, 66 MHz, 3.3 volt slot  ent1
U0.1-P1-I4   PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I5   PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I6   PCI 32 bit, 33 MHz, 5 volt slot    Empty
U0.1-P1-I7   PCI 32 bit, 33 MHz, 5 volt slot    Empty
U0.1-P1-I8   PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I9   PCI 64 bit, 66 MHz, 3.3 volt slot  ssa0
U0.1-P1-I10  PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I11  PCI 32 bit, 33 MHz, 5 volt slot    Empty
U0.1-P1-I12  PCI 32 bit, 33 MHz, 5 volt slot    scsi2
U0.1-P1-I13  PCI 64 bit, 66 MHz, 3.3 volt slot  Empty
U0.1-P1-I14  PCI 64 bit, 66 MHz, 3.3 volt slot  Empty

# drslot -c pci -Ia -s U0.1-P1-I5

The visual indicator for the specified PCI slot has
been set to the identify state. Press Enter to continue
or enter x to exit.

[Enter]

The visual indicator for the specified PCI slot has
been set to the action state. Insert the PCI card
into the identified slot, connect any devices to be
configured and press Enter to continue. Enter x to exit.

[Enter]
```

Figure 6. Example: Adding a Hot-Plug PCI Adapter

Remote I/O Ports and Additional I/O Drawers

RIO connections provide the dataflow between the CEC and the I/O drawers. RIO connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. Model 6M1 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.

The Model 6M1 CEC contains four RIO ports (A0/A1 and B0/B1) for creating two loops capable of supporting up to four I/O drawers. A one or two I/O drawer Model 6M1 system uses one RIO loop with ports A0 and A1. If a third and a fourth I/O drawer are attached to system, a second loop, with ports B0 and B1, must be created. Each port operates at 500 MHz in bidirectional mode and is capable of passing up to eight bits of data in each direction on each cycle of the RIO port. Therefore, the maximum data rate is 1 GB/s per I/O drawer.

The primary I/O drawer must be attached to the RIO-0 port of the CEC for system startup and access to the system through the support processor.

When adding secondary I/O drawers, each provides an additional 14 hot-plug capable PCI slots and two additional media bays (no additional connectors for external devices).

Cabling

Four types of interface cables are included with each system ordered:

- JTAG[6] (one per system)

  The JTAG cable connects the CEC with the primary I/O drawer. It is used for system control and initialization. This cable is available in two lengths: 3 m (# 5992) and 6 m (# 5993).

- V/S COMM cable (CEC SPCN – one per system)

  This cable also connects the CEC with the primary I/O drawer. It is used for power control and is also available in two lengths: 3 m (# 6132) and 6 m (# 6136).

- RIO Cables (two, three, five, or six per system)

  Remote I/O cables connect the CEC with the I/O drawers. The number of RIO cables that are necessary for a certain number of drawers are listed in Table 2. RIO cables can be ordered in three lengths (3 m # 3142, 6 m # 3143, and 15 m # 3144). For specific cabling information, refer to the *IBM @server pSeries 660 Model 6M1 Service Guide*, SA38-2538.

Table 2. Number of RIO cables needed

Number of drawers	Number of RIO cables	A loop	B loop
Only primary	2	X	-
One secondary	3	X	-
Two secondary	5	X	X
Three secondary	6	X	X

- SPCN[7] cables (zero to four per system)

  SPCN cables are used for power control of the secondary I/O drawers. These cables connect the primary with the secondary I/O drawers. When only a primary I/O drawer is used, no SPCN cables are necessary. With one secondary I/O drawer, two of these cables are necessary. For every additional secondary I/O drawer, another cable is necessary. For exact cabling information, refer to the *IBM @server pSeries 660 Model 6M1 Service Guide*, SA38-2538. These cables can also be ordered in three lengths (3 m # 6006, 6 m # 6008, and 15 m # 6007).

Dual Boot Bay Option

For booting from disks, external disk storage is required for the Model 6M1, because data storage internal to the drawers is not provided; however, optional boot bays are available in the primary I/O drawer. One or two SCSI IPL disks (non hot-plug) can be mounted with the IPL disk mounting hardware feature (# 6540). The internal IPL disk mounting hardware is available for use in the primary I/O drawer only and eliminates the use of PCI slots 13 and 14. Both disks can be connected to the internal SCSI adapter. If desired, one or both of the internal

[6] Joint Test Action Group (JTAG)
[7] Serial Power Control Network (SPCN)

disks can be connected to an external SCSI adapter port using the external SCSI adapter to internal IPL disk bays cable assembly (# 3139) and appropriate external SCSI cabling. The external SCSI adapter to internal IPL disk bays cable assembly (# 3139) provides a 68-pin P-style connector on the rear bulkhead of the I/O drawer. To attach the connector on the rear bulkhead to a VHDCI connector of a SCSI adapter, a 0.3 m (# 2118) converter cable, VHDCI TO P, mini 68-pin to 68-pin, is needed. If necessary, a SCSI cable 68-pin (P-type) to 68-pin (P-type) can be put in between (0.6 m # 2424, 2.5 m # 2425).

Figure 7 shows how the dual boot bay fits into the I/O drawer.

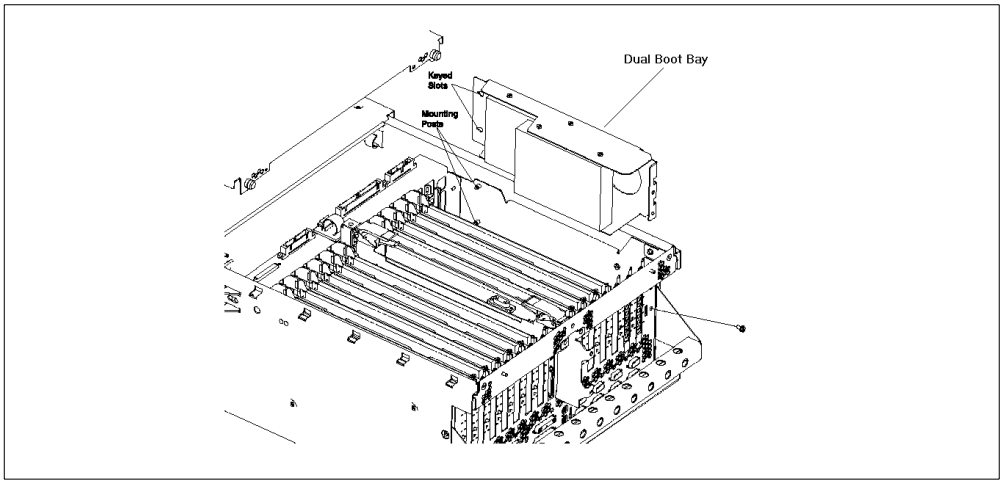

Figure 7. Dual Disk Bay Option

Software Requirements

The Model 6M1 requires AIX 4.3.3, with 4330-09 recommended maintenance package or later, or AIX 5L for POWER Version 5.1, with 5100-01 maintenance package or later.

In order to install the Model 6M1 from CD, you need an AIX 4.3.3 CD dated 04/2000 (LCD4-0286-05) or later, because the system will not boot from older AIX 4.3.3 CDs.

You can also download the actual maintenance level from the Internet to install the machine using NIM[8]. The URL to obtain an AIX 4.3 maintenance level is:

`http://techsupport.services.ibm.com/rs6k/fixes.html`

Instructions on how to obtain AIX 5L service is available at this site.

If you have problems downloading the latest maintenance level, ask your IBM Business Partner or IBM Representative.

[8] Network Installation Management (NIM)

Availability, Investment Protection, and Expansion

The following sections discuss how configurations, upgrades, and design features help you lower your cost of ownership.

High Availability Solution

Reliability of the system is further hardened by using the HACMP clustering solution available across the entire range of pSeries and RS/6000 servers. The HACMP solution exploits redundancy between server resources and provides application uptime. The Model 6M1 is available in a high-availability cluster solution package named the HA-6M1. This solution consists of the following components:

- Two pSeries Model 6M1 servers
- AIX Version 4.3.3 operating system (unlimited user license), or AIX 5L Version 5.1 (unlimited user license)
- HACMP 4.4.0 cluster software (APAR IY17684 required for AIX 5L), or later
- Two 7133-D40 SSA disk subsystem, with at least four disk drives each
- One 7014-T00 system rack
- All necessary redundant hardware and cables
- Serial ports S3 and S4 support HACMP heartbeat

This solution is sold at a price less than the sum of its parts. Ask your IBM Business Partner or IBM representative for further information.

Reliability, Availability, and Serviceability (RAS) Features

RAS features, such as redundant power supplies or N+1 hot-plug fans, have previously been discussed. Additional RAS topics are covered in this section.

Error Recovery for Caches and Memory

The RS64 IV processor L1 cache, the L2 cache, and the memory are protected by error correction code (ECC) logic. The ECC provides single-bit error correction and double-bit error detection for the L2 cache and the memory. All recovered error events are reported by an attention interrupt to the service processor, where they are monitored for threshold conditions.

The standard memory card has single error-correct and double error-detect ECC circuitry to correct single-bit memory failures. The double-bit detection helps maintain data integrity by detecting and reporting multiple errors beyond what the ECC circuitry can correct. In many cases (using DIMMs with 18 DRAMs, for example), memory chips are organized such that the failure of any specific memory module only affects a single bit within an ECC word (bit scattering), thus allowing for error correction and continued operation in the presence of a complete chip failure (chip kill recovery).

Another function, named *memory scrubbing*, provides a built-in hardware function, which performs continuous background reads of data from memory, checking for correctable errors. Correctable errors are corrected and rewritten to memory, and a threshold counter is maintained that will signal the service processor with special attention when the threshold is exceeded.

Chipkill™

The Model 6M1 provides Chipkill memory with selected memory options. Chipkill memory protects the server from any single memory chip failure and multibit errors from any portion of a single memory chip. Memory chip failures can cause server system crashes that can result in the permanent loss of business data. Chipkill DIMMs for the Model 6M1 provide the self-contained capability to correct real-time, multibit DRAM errors, including complete DRAM failures. This technology provides enhanced multibit error detection and correction that is not apparent to the system.

Note

The Chipkill memory technology is not supported on the # 4110 and # 4133 memory features.

Dynamic Processor Deallocation

The processors are continuously monitored for errors such as L2 cache ECC errors. When a predefined error threshold is met, an error log with warning severity and threshold exceeded status is returned to AIX. At the same time, the service processor marks the CPU for deconfiguration at the next boot. In the meantime, AIX will attempt to migrate all resources associated with that processor (tasks, interrupts, etc.) to another processor, and then stop the failing processor.

The capability of Dynamic Processor Deallocation is only active in systems with more than two processors, because device drivers and kernel extensions, which are common to multi-processor and uniprocessor systems would change their mode to uniprocessor mode with unpredictable results.

Dynamic Processor Deallocation is not supported when HMT is enabled.

Persistent Processor and Memory Deconfiguration

Processor and memory modules with a failure history are marked *bad* to prevent them from being configured on subsequent boots. This history is kept in the VPD[9] records on the FRU[10], so the information moves physically with the FRU and is cleared when the FRU is replaced and stays with the failed FRU when it is returned to IBM. A CPU or memory module is marked bad when:

- It fails BIST[11]/POST[12] during boot (as determined by CEC initialization code in the service processor)
- It causes a machine check or check stop during runtime and the failure can be isolated specifically to that CPU or memory module (as determined by the service processor)
- It reaches a threshold of recovered failures (for example, ECC correctable L2 cache errors, see preceding) that result in a predictive call-out (as determined by service processor)

[9] Vital Product Data (VPD)
[10] Field Replaceable Unit (FRU)
[11] Built-In Self-Test (BIST)
[12] Power-On Self-Test (POST)

During CEC initialization, the service processor checks the VPD values and does not configure CPUs or memory that are marked bad, much in the same way that it would deconfigure them for BIST/POST failures.

I/O Expansion (RIO) Recovery

The RIO interface supports packet retry on its interface, which means that it will automatically try to resend a packet if it gets no acknowledgment or a bad response until a time-out threshold is reached.

RIO also supports a closed-loop topology configuration, which is required for pSeries and RS/6000 products. A pair of RIO ports are each connected in daisy-chain fashion as primary ports to one or more RIO nodes, with the far ends of the daisy-chains interconnected to close the loop. RIO hubs will automatically attempt to reroute packets through the alternate RIO port if a successful transmission cannot be completed (for example, the retry threshold is exceeded) through the primary port. In this way, no single link failure in the RIO loop will cause any impact to system operation, although the failure will be reported for deferred maintenance.

PCI Bus Error Recovery

As described in the PCI slot section, every slot is connected through a PCI-to-PCI bridge chip to a primary PCI bus. Therefore, each slot is logically and physically isolated onto its own individual PCI bus. This fact provides a special error handling mode that allows the bridge chip to *freeze* access to an adapter when a PCI bus error occurs on the interface between that adapter and bridge chip. In this frozen mode, DMAs[13] are blocked, stores to that device address space are discarded, and loads result in a return value of all 1s. Device drivers can be programmed to look for these dummy responses on loads and can attempt recovery. The AIX 5L Version 5.1 is required for this support.

System Power Control Network (SPCN)

SPCN consists of a set of power/environmental controllers, one in each system drawer, interconnected by a set of serial communication links (see “Cabling” on page 16). The SPCN node in the same drawer as the service processor is the master or primary SPCN, while all others are slaves or secondary. SPCN provides the following functions:

- Powering all the system drawers up or down, when requested.

  The SPCN hardware has connections to the VPD that is resident on each of the pluggable cards and the backplane. The VPD is located on each of the cards in the form of an I^2C chip. This chip is accessed during initial power on sequence and the data contents are read by the service processor. Using this function the service processor decides not to use components that are marked *bad.*

- Powering down all the system drawers on critical power faults.
- Monitors power, fans, and thermal conditions in the system for problem conditions, which result in an EPOW. EPOW stands for environmental and power off warnings and is a function to inform the service processor or the operating system early about an event that happened in the hardware. There are different warnings, such as cooling warnings or power fail warnings, which result in entries in the error log. If there is a serious error, such as the temperature reaches a specific limit, the system will be shutdown.

[13] Direct Memory Access (DMA)

- Reporting power and environmental faults, as well as faults in the SPCN network itself, on operator panels and through the service processor.
- Assigning and writing location information into various VPD elements in the system.

Service Processor

The Model 6M1 has an integrated enhanced service processor, located in the primary I/O drawer. When the system is powered down, but still plugged into an active power source, the service processor functions are still active under standby power. This function provides enhanced RAS by not requiring AIX to be operational for interfacing with a system administrator or service director for RS/6000. All service processor menu functions (using the local, remote, or terminal concentrator console), as well as dial out capability, are available even if the system is powered down or unable to power up. The next sections describe selected features of the enhanced service processor.

Automatic Reboot

The system will automatically reboot (if the appropriate policy flags are set) in the following conditions:

- Power is restored after a power loss during normal system operation.
- Hardware Checkstop Failures.
- Machine Check Interrupt.
- Operating System Hang (Surveillance Failure).
- Operating System Failure.

Surveillance

The service processor, if enabled through service processor setup parameters, performs a surveillance of AIX through a heartbeat mechanism. If there is no heartbeat within the time-out period, the service processor does the following:

- Creates a system reset to allow an AIX dump to occur.
- Upon receiving a reboot request (either after the dump, or immediately, if dump is not enabled), the service processor captures scan debug data for the system.
- Reboots the system.

Dial-Out (Call Home), Dial-In

If enabled, the service processor can dial a pre-programmed telephone number to report errors. If enabled, it is also possible to access the service processor remotely through a modem connection. When the service processor is in standby mode because the system is powered off or an error occurred, the service processor monitors an incoming phone line to answer calls, prompts for a password, verifies the password, and remotely displays the standby menu. The remote session can be mirrored on the local ASCII console if the server is so equipped and the user enables this function.

Processor and Memory Boot Time Deconfiguration

As described previously, processors can be dynamically deconfigured by the system. It is also possible to deconfigure processors and also memory with menus of the service processor for benchmarking reasons. For further

information, refer to the *IBM @server pSeries 660 Model 6M1 Service Guide*, SA38-2538.

Note

If the memory is to be temporary deconfigured (for benchmarking or sizing, for example), it is also possible to use the AIX `rmss` command to simulate a specific amount of memory (only below the real memory limit).

Fast Boot
This feature, set as default, allows you to select the IPL type, mode, and speed for your boot capabilities using service processor menus. Selecting fast boot results in several diagnostic tests being skipped and a shorter memory test being run. Therefore, the startup process is faster, but possible problems might not be discovered at startup.

Service Processor Restart
The service processor design for the Model 6M1 includes the ability to reset the service processor. This enables the system firmware to force a hard reset of the service processor if it detects a loss of communication. Since this would typically occur while the system is already up and running, the service processor reset will be accomplished without impacting system operation.

Boot to SMS Menu
The boot mode menu allows someone to select other things to boot to the SMS menu. This function provides booting into SMS menu without pressing a key. This function is useful because it is not necessary to wait in front of the system and press F1 (graphic display) or 1 (ASCII terminal) at the right moment.

Capacity Upgrade on Demand

Capacity Upgrade on Demand (CUoD) is a new feature that allows you to have inactive processors installed on your system, which can be made active quickly and as easily as your business needs require. When more processing capacity is required, you issue the AIX `chcod` command to increase the number of processors, in increments of two, from four to eight. The additional processors are activated on the next reboot.

Capacity Upgrade on Demand is available only to end user customers who sign contracts for the service. The CUoD contracts are available from your IBM or IBM Business Partner representative.

Capacity Upgrade on Demand for Model 6M1 is available only for the second processor slot. The first processor slot must be occupied by a 4-way, 750 MHz processor card (# 5213). Systems with feature # 5213 in the first processor slot may be configured with a second 4-way processor card in the second slot, while leaving two or four of the processors inactive. Feature # 8306 is a 4-way processor card with no processors activated. Feature # 8307 is a 4-way processor card with two processors activated.

AIX software preinstall (# 5005) is required on all initial system orders incorporating Capacity Upgrade on Demand processor features. The Model 6M1 system periodically checks the number of processors in use, detecting that the customer has increased the number of active processors in the system. When the

change is detected, the Model 6M1 notifies IBM that additional processing power is being utilized.

The customer is responsible for placing an order for the additionally activated processors with their IBM or IBM Business Partner sales representative.

IBM Electronic Service Agent must be installed and active on all systems containing CUoD 0-way active (# 8306) or 2-way active (# 8307) processor features. There is no longer a prerequisite for eligible CUoD machines to be under IBM warranty or IBM Service for Machines.

IBM requires that you install and run the IBM Electronic Service Agent for each eligible pSeries or RS/6000 machine in accordance with IBMs documentation.

Unused CUoD processors can also be utilized to replace system processors which have been deallocated due to failure. A system reboot will activate one of the unused CUoD processors to replace the failing processor until a replacement can be installed.

Unactivated processors cannot be tested by AIX diagnostics, nor do they appear when using AIX commands or diagnostics. However, Systems Management Services and Service Processor menus are not affected by CUoD, and all processors are tested at boot time.

System Upgrades

For owners of a 7026-6H1, it is possible to upgrade to Model 6M1 while keeping the original serial number.

Model 6H1 systems converted to Model 6M1 systems will require replacement of their system CEC. Model 6H1 4-way or 6-way processor cards can be replaced with Model 6M1 4-way processor cards via feature conversion. This processor conversion is available at the time of the initial model upgrade only. A maximum of one 6H1 processor card can be converted to one Model 6M1 processor card for each system being converted. The existing Model 6H1 processor card being replaced is returned to IBM along with the CEC.

The memory DIMMs can move from Model 6H1 to Model 6M1, but it should be considered that DIMMs in Model 6M1 have to be installed in octals.

Because the I/O drawers in Model 6H1 are the same as in Model 6M1, they can move to Model 6M1, but some adapters that are supported in Model 6H1 are not in Model 6M1. For further information about the support of adapters, refer to the *PCI Adapter Placement Reference Guide*, SA38-0583.

All racks used for Model 6H1 can also be used with Model 6M1.

For owners of a 7026-M80, or Model 6M1 with 500 MHz processors, upgrades to 750 MHz processor cards require a clock card and power regulator, which are included with the upgrade.

For owners of a 7026-M80, upgrades to 750 MHz processors also require a new CEC backplane.

Check your IBM representative for the current availability of these upgrades.

External Storage Expandability

The storage requirement for the Model 6M1 is met externally through several IBM storage options. The storage subsystems can be connected externally as a standalone tower or from within another rack. Disk storage capacity can also be provided externally through storage servers. Using differential Ultra SCSI, the Model 6M1 can be attached to the IBM Enterprise Storage Server. By using the Fibre Channel Adapter, the Model 6M1 can be attached to the IBM Fibre Channel RAID Storage Server or the IBM Enterprise Storage Server. The IBM storage subsystems, which also can be mounted in the available space of the same rack as the Model 6M1, include the following:

2104-DU3 The IBM 2104 Model DU3 is a drawer (rack mount) model with 14 disk drive bays, 509.6 GB maximum capacity and up to two Ultra3 SCSI interface/ports. It requires 3 EIA of rack space.

7133-D40 The IBM 7133 is a serial storage architecture disk system. It provides 4 to 16 advanced disk drive modules. Combinations of 9.1 GB, 18.2 GB, and 36.4 GB SSA drives can provide capacity points ranging from 36 GB to over 582 GB. It requires 4 EIA of rack space.

7337-306 The IBM 7337-306 is a digital linear tape library and it contains one or two DLT tape drives, providing a data storage capacity of 525 GB/1050 GB in uncompressed/compressed format. It requires 5 EIA of rack space.

Within the rack, IBM Magstar tape subsystems can be mounted. These include 3590-E11 (two requires 12 EIA), 3590-B11 (two requires 12 EIA), 3570-C12 (6 EIA), 3570-C11 (6 EIA), 3490-F11 (4 EIA).

SP Attachment and Clustering

The joining of the Model 6M1 to the RS/6000 SP™ satisfies the need that many SP environments have for large, powerful, and memory rich processors for their database servers and today's e-business applications. All SP-attached systems have an optional connection to the SP Switch. Thus, they can utilize the SP advantage of high-speed interconnect.

Clustering, in a computer context, has became popular during the past five years for a number of reasons:

- Clusters offer scalability, which is essential to resolving complex computer problems.
- Consolidation with a single point of control to reduce management costs.
- Improved availability of resources through sharing, replication, and redundancy within a cluster.

A maximum of 32 Model 6M1 servers are supported in one Clustered Enterprise Servers (CES) system. This configuration does not contain a physical SP frame or SP Switches. It is a clustered set of Model 6M1 servers centrally managed through Parallel System Support Programs (PSSP) software from a single hardware control point, the control workstation. A 9076 Model 555 provides SP Switch attach without the requirement to order an SP node.

For further information about SP, see the following Web site:
`www.rs6000.ibm.com/resource/aix_resource/sp_books/index.html`

Reference

The following sections list additional materials available for further research.

System Documentation

For more detailed information, refer to the following documents:

- *IBM @server pSeries 660 Model 6M1 Installation Guide*, SA38-0576
- *IBM @server pSeries 660 Model 6M1 User's Guide*, SA38-2537
- *IBM @server pSeries 660 Model 6M1 Service Guide*, SA38-2538
- *PCI Adapter Placement Reference Guide*, SA38-0538

Select IBM Redbooks

The following IBM Redbooks are related to the material discussed in this paper:

- *RS/6000 Systems Handbook 2000*, SG24-5120
- *IBM @server pSeries 680 Handbook Including RS/6000 Model S80,* SG24-6023
- *IBM Enterprise Storage Server,* SG24-5465
- *Monitoring and Managing IBM SSA Disk Subsystems,* SG24-5251
- *AIX 5L Differences Guide Version 5.1 Edition,* SG24-5765
- *NIM: From A to Z in AIX 4.3,* SG24-5524
- *AIX Logical Volume Manager, from A to Z: Introduction and Concepts,* SG24-5432
- *Understanding IBM @server pSeries Performance and Sizing*, SG24-4810

Select Internet Links

For more detailed information see the following Web sites:

```
www.ibm.com/servers/eserver/pseries/
www.ibm.com/servers/eserver/pseries/library/hardware_docs/
www.rs6000.ibm.com/support/micro/
www.ibm.com/servers/aix/
www.chips.ibm.com/
www.storage.ibm.com/
www.hursley.ibm.com/~ssa/rs6k/
www.redbooks.ibm.com/
```

Acknowledgements

Assistance creating this white paper came from the following individuals and was appreciated. Thank you!

Atkins, Stephen	IBM U.K.
Amy, Larry	IBM Austin
Behr, Lisa	IBM Bethesda
Frey, Brad	IBM Austin
Gray, Forrest	IBM Austin

Grzinich, John	IBM Austin
Haug, Volker	IBM Germany
Henderson, Daniel	IBM Austin
Jackson, Tenley	IBM Southfield
Lee, Dennis	IBM Austin
McCord, Scott	IBM Austin
Mihaltse, Bill	IBM Austin
Ruth, Dave	IBM Austin
Scherrer, Carolyn	IBM Austin
Shempert, Craig	IBM Austin
Stys, Mike	IBM Austin
Toutant, Ed	IBM Austin
Vicknair, Brian	IBM Austin

Biographies

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

Shyam Manohar is the Product Manager with the pSeries ASEAN team. He has over nine years selling experience in the RISC UNIX marketplace. Shyam has contributed over five years of selling for the RS/6000 and pSeries brands. He holds unique expertise in product positioning in competitive scenarios. Shyam is also the sales evangelist for high-end servers in the ASEAN region.

Special Notices

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 IBMs product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBMs intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

This document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including

these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.