# IBM @server pSeries 660 Model 6H1

## Technical Overview and Introduction

April 17, 2001

Gregor Linzmeier


Scott Vetter

International Technical Support Organization

Austin, TX

IBM

# IBM @server pSeries 660 Model 6H1

The pSeries product line is part of the IBM @server brand - a new generation of servers featuring unmatched availability and scaleability, broad support of open standards for the development of portable new applications, and IBM @server Advantage for managing the unprecedented demands of e-business.

On April 17, 2001, IBM introduced the IBM @server pSeries 620 Model 6F1 and IBM @server pSeries 660 Model 6H1 to enhance the mid-range server lineup with products using many of the design elements that led to the success of the high-end pSeries Model 680.

## Overview

While the pSeries 620 Model 6F1 and pSeries 660 Model 6H1 are the systems that provide a growth path for existing installations of IBM/RS/6000 Model F80s and H80s, the Model M80 is designed to provide leadership performance among the mid-range 8-way systems.

The Model 6H1 offers a rack mounted alternative to Model 6F1 with the flexibility to select an external storage solution or integrate into an existing one, to closely match the storage needs.

This paper discusses, in detail, the processor, memory, I/O, expandability, reliability, and other technical aspects related to Model 6H1.

An update to this paper, indicated by the change bars, was made to reflect the new processor features available after April 9 2002.

# pSeries 660 Model 6H1 Introduction

Model 6H1 is a member of the 64-bit family of symmetric multiprocessing (SMP) enterprise servers from IBM.

The Model 6H1 provides scalability using the 64-bit RS64 IV processor packaged in 1-, 2-, 4-, or 6-way 750 MHz, in 1-, 2-, or 4-way 600 MHz, 6-way 668 MHz, or 1-, 2-, or 4-way 450 MHz processor configurations with a maximum of 32 GB of real memory. Model 6H1 incorporates an I/O subsystem supporting 32-bit and 64-bit standard PCI adapters. The architecture allows 32- and 64-bit applications to run simultaneously.

The following sections discuss the key components and their packaging with regards to the Model 6H1.

## Central Electronics Complex and I/O Drawer

The basic Model 6H1 package consists of a Central Electronics Complex (CEC) and an I/O drawer. The CEC drawer incorporates the system processors, memory, and supporting system logic. The CEC is connected with different cables to the primary I/O drawer (see “Cabling” on page 16).

The primary I/O drawer provides 14 hot-plug PCI slots, the connectors for external devices, and the service processor. If more than 14 slots are necessary, it is possible to add a secondary I/O Drawer (see “Considerations for a Second I/O Drawer” on page 16), thus providing an additional 14 hot-plug PCI slots.

Both the CEC and the I/O drawer are 5 U (EIA Units)[1] in height. The exact size of both the CEC and I/O drawer is 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 the CEC or the I/O drawer is from 41 kg (90 lbs) to 52 kg (115 lbs), depending on the configuration.

The Model 6H1 is designed to be installed and maintained by trained service representatives.

The drawers can be mounted in an 7014 Model S00 or in the 7014 Model T00 and 7014 Model T42 racks (see “Enterprise Racks” on page 3).

Figure 1 shows the front view of a Model 6H1 with two I/O drawers and a Model 7316-TF1 Flat Panel Console in a rack.

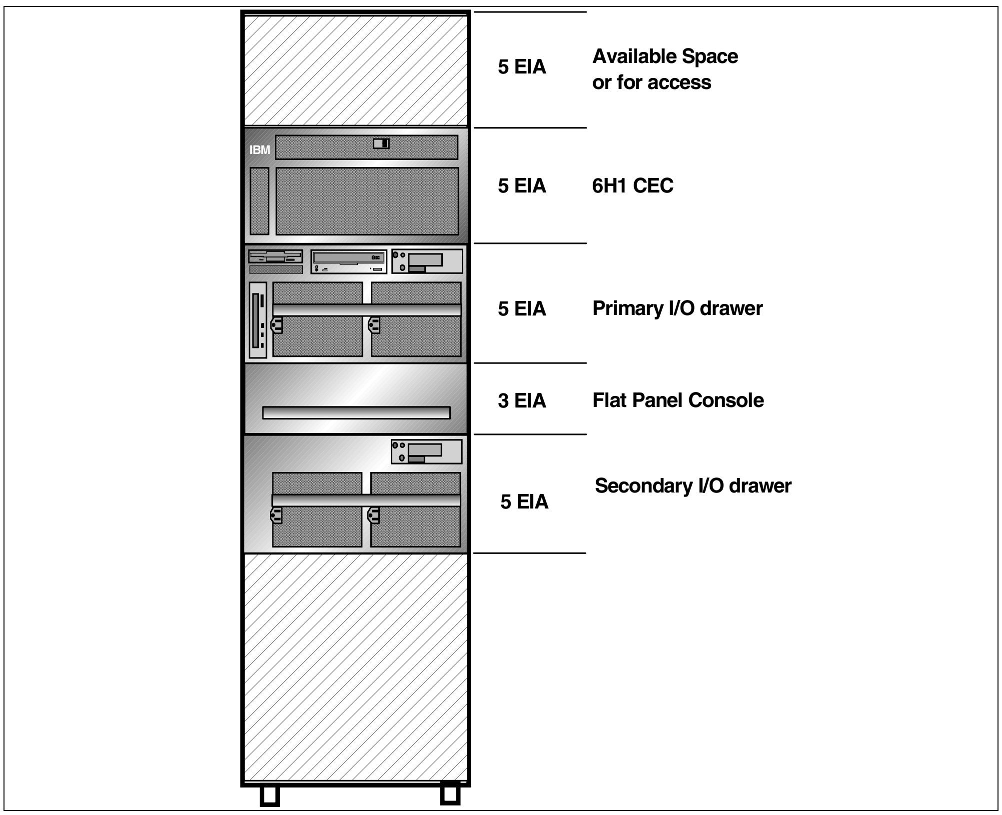


*Figure 1. Model 6H1 with Two I/O Drawers and Flat Panel Console in a T00 Rack*

## 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). 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). In a standard configuration, external disk storage is required for the Model 6H1 as data storage; however, optional

[1] One EIA Unit (Electronic Industries Association Unit) is 44.45 mm or 1.75’’

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 standard configuration, both the CEC drawer and the I/O drawer are equipped with one power supply in the right power supply bay. These power supplies can be either AC (# 9172) or -48V DC (# 9175), and both are designed to support up to 6-way SMP systems with 32 GB of memory. For both drawers, redundant AC power supplies with a separate power cord are available (# 6282 for the CEC and # 6283 for the I/O drawer). The power supplies of the CEC and the I/O drawer cannot be interchanged. For DC power, a redundant power supply is standard, and no additional power features are required. All power supplies are hot-pluggable and allow concurrent repair when redundant power is installed.

The Model 6H1's CEC and I/O drawer cooling subsystems have N+1 hot-pluggable fans, which are performance monitored by SPCN (see “System Power Control Network (SPCN)” on page 21). In the event of a single fan failure, the other fans increase their speed to provide sufficient cooling. The hot-pluggable fans can also be replaced concurrently if the optional redundant power feature is installed.

## Enterprise Racks

The following discussion is provided as an overview of racks 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. The rack provides increased capacity, greater flexibility, and improved floor space utilization.

### 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 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 (1795 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 tall enclosure to house the maximum amount of equipment in the smallest possible floor space. The features that differ in the Model T42 rack from the Model T00 include the following:

- 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 (2045 pounds)

### Rack Mounting Rules for Model 6H1

There are some rules that should be considered when mounting drawers into a rack. These are as follows:

- The Model 6H1 I/O drawers may be placed at any location in the rack.
- If mounting a Model 6H1 CEC above the I/O drawer, a minimum of 5 U is required between the CEC and the top of the rack for service because the cover of the CEC interferes with the top of the rack when it is opened. If the CEC is located under another drawer, the cover can be opened because of the open space within the rack offered behind the drawer above.
- The space between the top of the rack and the CEC can be used for storage, such as 7133-D40 or an I/O drawer. These devices do not interfere with the top of the drawer beneath because it does not extend the full depth of the rack.

A Model 6H1 with one I/O drawer is 10 U in height, so a maximum of three Model 6H1s fit in an S00 rack or T00 rack, or a maximum of four in a T42 rack.

### IBM 7316-TF1 Flat Panel Display 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 T540 features a 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 console has the following attributes:

- 3 EIA units (3 U)
- IBM T540 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 GXT130P Graphics Adapter (#2830) or follow-on graphics adapter.

## 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 6H1 CEC. In this section, the different physical components and the SMP processor configurations in the system schematic are discussed.

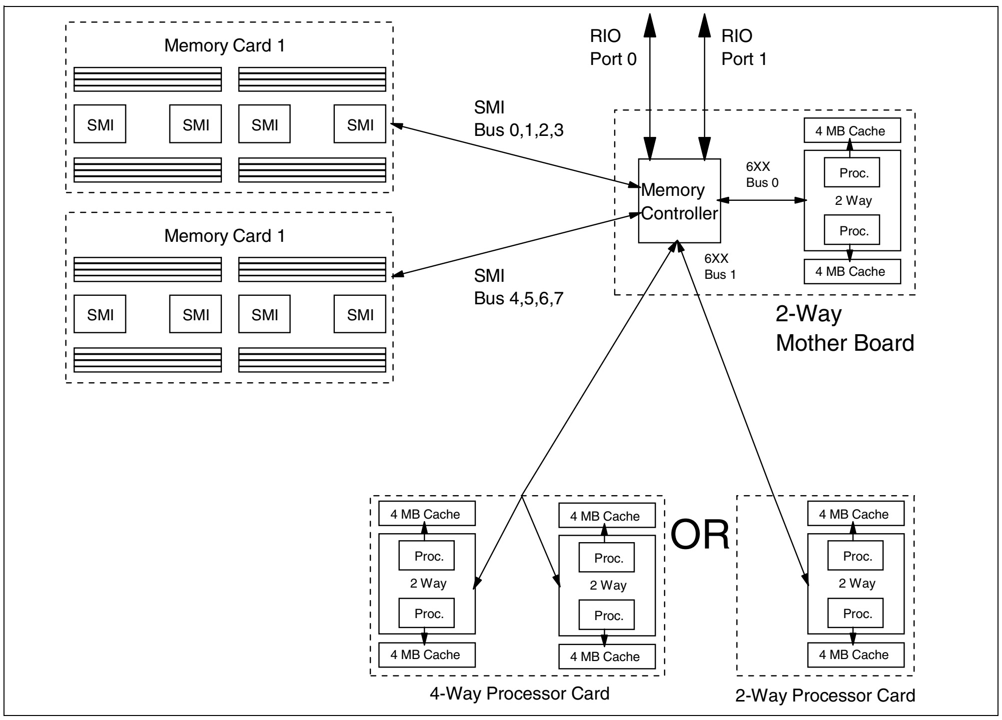


*Figure 2. Model 6H1 System Schematic for 2-, 4-, or 6-Way 600 or 668 MHz SMP within the CEC*

#### RS64 IV RISC Processor

The RS64 IV processor card used in the 6H1 has the following attributes:

- 600 MHz, 668 MHz, or 750 MHz operating frequency
- 128 KB on-chip L1 instruction 2-way associative with parity and retry for data integrity
- 128 KB on-chip L1 data cache with ECC, 2-way associative
- On-chip L2 cache directory
- 8 MB of off-chip L2 cache using ECC double data rate (DDR) SRAM per processor for 1-, 2-, 4-, 6-way at 750 MHz, 8 MB of off-chip L2 cache using ECC double data rate (DDR) SRAM per processor for 6-way SMP at 668 MHz,

4 MB of off-chip L2 cache using ECC double data rate (DDR) SRAM per processor for 2- and 4-way SMP at 600, and 2 MB of off-chip L2 cache using ECC Single Data Rate (SDR) SRAM for a 1-way system at 600 MHz.

- PowerPC 6xx bus architecture, 16-byte wide bus interface

The RS64 IV processor is available in three operating frequencies: 600 MHz, 668 MHz, and 750 Mhz. The frequency is accomplished by leveraging IBMs copper (CMOS 8S) 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.

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 by overlapping memory access with other computation. AIX Version 4.3.3 and AIX 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. HMT does not work on single processor systems. For further information, see /usr/lpp/bos/README.HMT as shipped with the AIX operating system.

### Processor Boards

The processor boards used in the Model 6H1 for 1-, 2-, 4-, and 6-way SMP configurations come in the form of a single book, and are described as follows:

**Single Processor 600 MHz**

As shown in Figure 3 on page 7, a single processor board consists of a single RS64 IV processor operating at 600 MHz, on-board memory slots, and a memory controller in a single book. Upgrades to additional processors require changing of the processor book. However, the single processor board is a cost-reduced package.

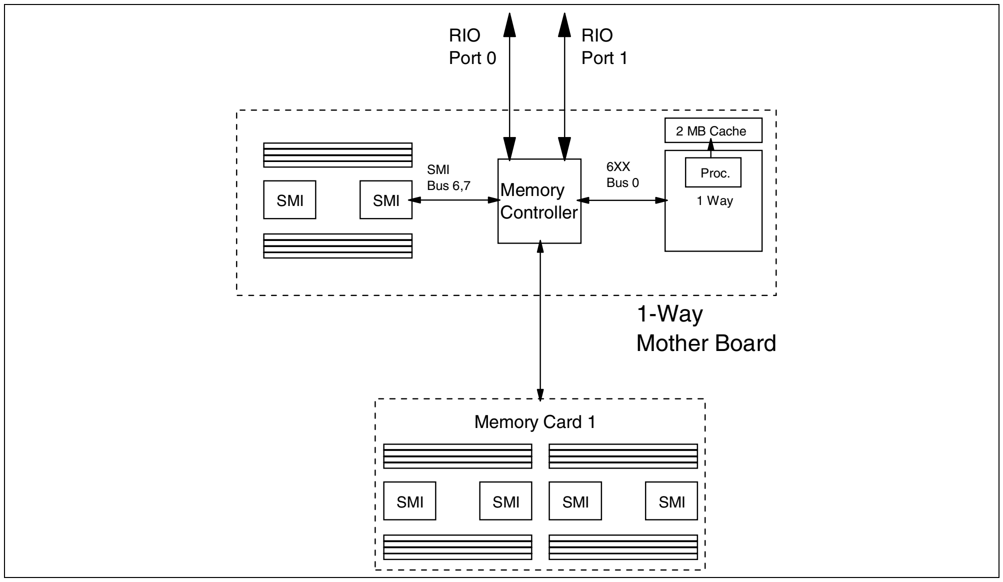


*Figure 3. Model 6H1 System Schematic for 1-Way Processor within the CEC*

**2- and 4-Way SMP 600 MHz**

A 2-way SMP configuration is provided by a processor board consisting of a pair of RS64 IV processors operating at 600 MHz and a memory controller. Expansion to 4-way SMP is provided by interfacing an additional processor board consisting of a pair of RS64 IV processors operating at 600 MHz. However, the upgrade from 2-way to 4-way SMP is offered as a book swap because the interface between the mother card and daughter card requires a factory tolerance that is unable to be reproduced in the field.

**6-Way SMP 668 MHZ**

A 6-way SMP configuration uses two processor boards interfaced to each other. One processor board consists of a pair of RS64 IV processors operating at 668 MHz and a memory controller. The other processor board consists of four RS64 IV processors operating at 668 MHz. Upgrades from a 2- or 4-way SMP to a 6-way SMP are offered as a book swap because the interface between the mother card and daughter card requires a factory tolerance that is unable to be reproduced in the field.

**1-, 2-, 4-, and 6-Way SMP 750 MHz**

The new processors are offered on 1-, 2-, 4-, and 6-way cards running at 750 MHz. The Level 2 cache per processor is 8 MB for all 750 MHz processor cards. A mew manufacturing process enables IBM to place processor modules on to a single motherboard for the 4-way and 6-way configurations as shown in Figure 4 on page 8. 1- and 2-way cards are manufactured as before. Elimination of the mother/daughter card connector dramatically increases the reliability.

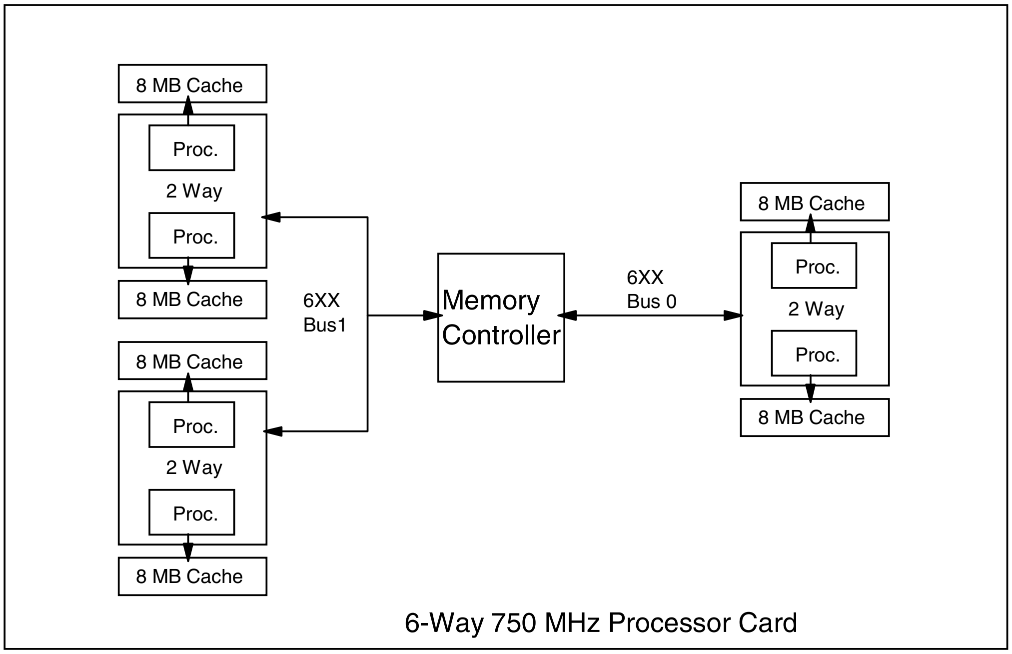


*Figure 4. pSeries 660 Model 6H1 System Schematic for 6-Way Processor 750 MHz*

### Memory Controller

A single custom chip provides the function of the memory controller and the I/O hub in the Model 6H1. The controller chip provides interfaces to processors, memory, and the I/O subsystem.

The RS64 IV processors on the processor boards are connected to the memory controller through the PowerPC 6xx bus. The controller chip is a part of the first processor board. The memory controller provides a single 6xx bus interface in a single processor configuration. For 2-way SMP configurations, the controller provides a 6xx bus interface to the pair of RS64 IV processors present in the same board. The memory controller provides another 6xx bus interface for CPU expansion using an additional processor board. The 4- and 6-way SMP configurations consists of a total of two processor boards that use the two 6xx bus interfaces provided by the memory controller installed in a book.

In the Model 6H1, the 6xx bus is a 16-byte wide bus and the operating clock rate of the bus depends on the processor clock speed. The 6xx bus operates at a clock rate of 150 MHz with a processor clock speed of 600 MHz or 750 MHz. And for a processor clock speed of 668 MHz, the 6xx bus operates at a clock rate of 133.6 MHz.

### Memory Subsystem

The memory controller provides two memory bus interfaces and provides the reliability functions of ECC as well as memory scrubbing. Memory scrubbing provides a built-in hardware function designed to perform continuous background reads of data from memory to check for correctable errors. The memory configuration for a single processor configuration and 2-, 4-, or 6-way configurations is:

- In a single processor configuration, the on-board memory, consisting of eight DIMM slots, is interfaced to one of the two memory interfaces in the controller.

The other interface is used by a separate riser memory card. The riser memory card provides 16 DIMM slots. While the DIMM slots in the on-board memory are populated in pairs, the slots in the riser memory card are populated in quads. The minimum configuration requires a pair of DIMMs in the on-board memory. Once the on-board memory slots are filled and more memory capacity is desired, the DIMMs are moved to the riser memory card and the next increment is made as a quad. The single processor configuration can provide a maximum memory of 16 GB by using the riser memory card and populating each of the 16 slots with 1024 MB DIMMs.

- In 2-, 4-, or 6-way SMP configurations, the memory is provided using two separate riser memory cards, each with 16 DIMM slots and populated with DIMMs in quads. The two riser cards are interfaced to the two memory interfaces on the memory controller. The minimum 2-way SMP configuration requires a single riser memory card populated with a quad of DIMMs. The second riser memory card with minimum of a quad of DIMMs can be configured only after the 16 DIMM slots in the first riser memory card are fully populated. 2-, 4-, or 6-way processor configuration can support up to 32 GB maximum memory by using the two memory riser cards fully populated with 1024 MB DIMMs.

The Model 6H1 uses 200-pin 10ns SDRAM DIMMs. DIMMs of equal sizes must be used while populating in pairs or quads. DIMM size used in one pair or quad can, however, coexist with a different DIMM size used in another pair or quad.

In the Model 6H1, the bus interface from each riser memory card to the memory controller is 8-bytes wide (two interfaces each 4-bytes wide) and operates at a clock rate of 300 MHz, or 267.2 MHz. No additional memory bandwidth can be achieved by splitting memory between cards.

**Bus Bandwidth**

The following are the theoretical maximum bandwidths as applicable for a 6-way, 668 MHz SMP configuration:

- Memory port bandwidth: 2.14 GB/s
- Processor port bandwidth: 2.14 GB/s
- Single drawer I/O bandwidth: 1 GB/s (500 MB/s bi-directional)
- Two drawer aggregate I/O bandwidth: 2 GB/s (2 x 500 MB/s bi-directional)

The following are the theoretical maximum bandwidths as applicable for 2- or 4-way, 600 MHz SMP configurations:

- Memory port bandwidth: 2.4 GB/s
- Processor port bandwidth: 2.4 GB/s
- Single Drawer I/O bandwidth: 1 GB/s (500 MB/s bi-directional)
- Two drawer aggregate I/O bandwidth: 2 GB/s (2 x 500 MB/s bi-directional)

The following are the theoretical maximum bandwidths as applicable for 1-, 2-, 4-, or 6-way, 750 MHz SMP configurations:

- Memory port bandwidth: 2.4 GB/s
- Processor port bandwidth: 2.4 GB/s
- Single Drawer I/O bandwidth: 1 GB/s (500 MB/s bi-directional)

- Two drawer aggregate I/O bandwidth: 2 GB/s (2 x 500 MB/s bi-directional)

**I/O Hub Function**

The controller also functions as the I/O hub. The controller provides two RIO (remote I/O) ports. The ports are connected to the I/O subsystem in the I/O drawer in a loop-like fashion. Each RIO port has two uni-directional, 1-byte wide links. All the I/O transfers take place using two primary RIO ports, which operate at 500 MHz (500 MB/s bi-directional or an aggregate of 1 GB/s).

## I/O Subsystem Architecture

The I/O subsystem consists of a primary I/O drawer and, if more PCI slots are necessary, of a secondary I/O drawer also. These I/O drawers in Model 6H1 are the same as the I/O drawers in the Model M80.

The primary I/O drawer contains the following features:

- Fourteen available hot-plug PCI slots
- Service processor (only in 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[2] 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
- 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)
- One parallel port (bi-directional)

The Debug port (3) is for diagnostics and is normally covered with a metal plate. It uses the same connector as the parallel port. To avoid confusion, this port should always remain covered.

SPCN connectors J11 and J14 should not be used.

Figure 5 shows the rear view of the primary I/O drawer and the port locations.

[2] Very High Density Cable Interconnect (VHDCI)

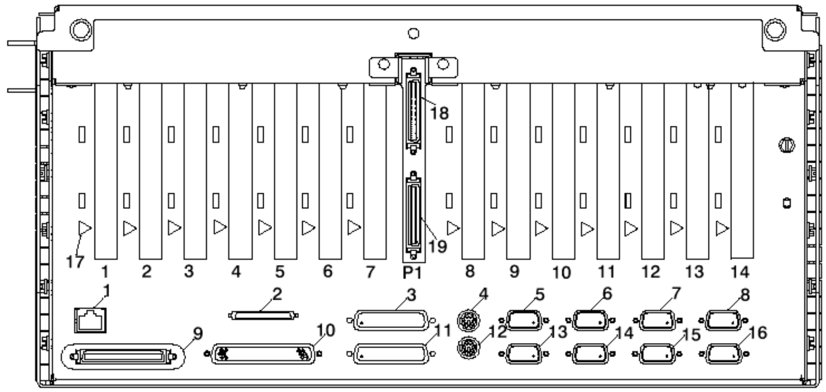

| | | | |
|---|---|---|---|
| 1 | Ethernet connector | 2 | SCSI connector |
| 3 | Debug | 4 | Mouse connector |
| 5 | Serial port 1 | 6 | Serial port 2 |
| 7 | J11 (SPCN 1 connector, reserved) | 8 | J14 (SPCN 3 connector, reserved) |
| 9 | V/S COMM | 10 | JTAG connector |
| 11 | Parallel port | 12 | 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 5. 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 6 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, and 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 by only having access to this button through a *pinhole* in the operator panel cover. This button is for service use only.

- Speaker (beeper)

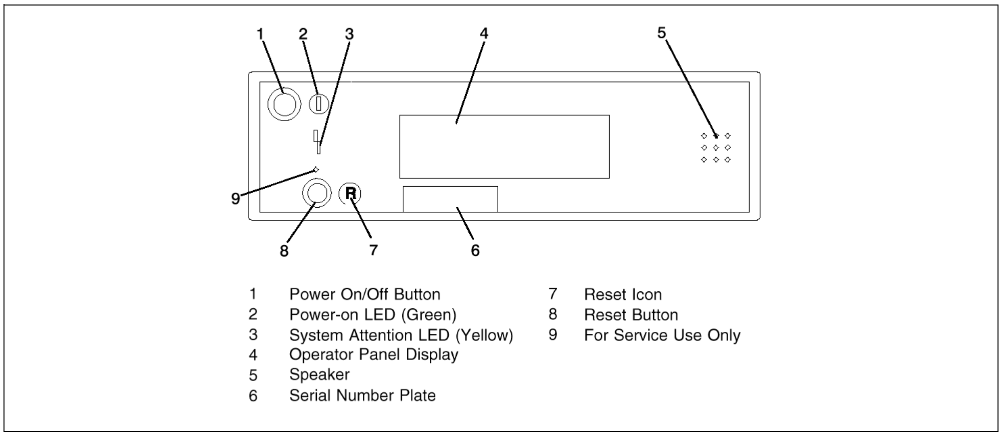


*Figure 6. Operator Panel*

### I/O Drawer Electronics

Figure 7 shows the major electronic components of the I/O drawer. They are discussed in the following section.

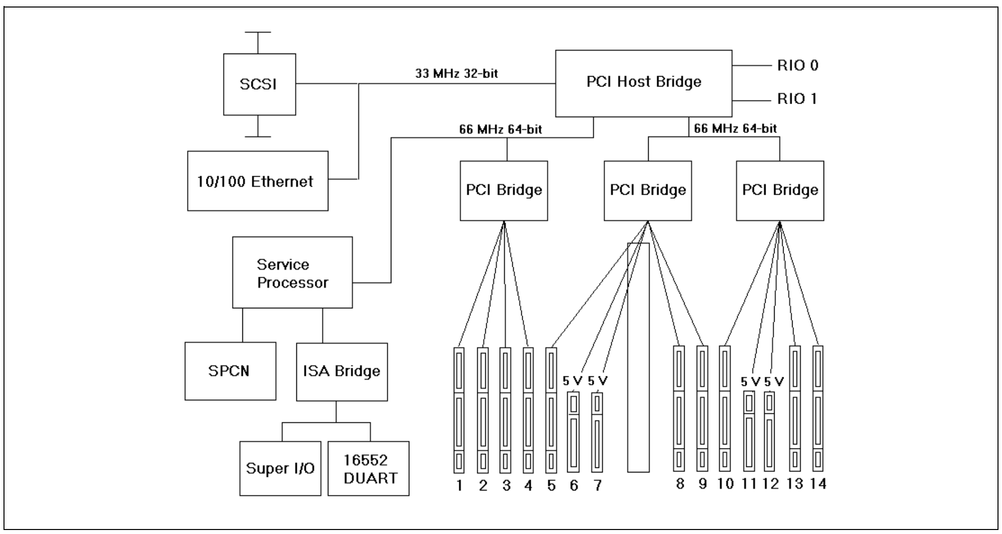


*Figure 7. Slot Layout of the I/O Drawer*

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. The reason for this architecture is to prevent damage to the components in the I/O drawer when plugging in hot-plug PCI cards while the system is running. Other features will 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 provides the two RIO ports to connect to the CEC, and on the other side, three primary PCI busses. 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, Ultra-2 SCSI external) and the on-board 10/100 Mb/s 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[3], 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 3.3 V signaling at 66 MHz in contrast to the four 32-bit slots that operate at 5 V signaling at 33 MHz (Figure 7). When adding adapters to the system, it is important at which signaling the adapter works: 3.3 V, 5 V, or universal, which means the adapter works at both voltages. That is, for example, 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 offers 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 and reduce performance potential.

#### Hot-Plug PCI Adapters

The function of hot-pluggable PCI adapters is to provide concurrent adding or removal of PCI adapters when the system is 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 5) 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.

[3] 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-pluggable, 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 which 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. There are three possibilities for managing hot-plug PCI slots in AIX:

- Command line:
  - `lsslot` - Lists 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 5 on page 11) 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 Identify state (LED flashes slowly) to verify the right slot was selected. After pressing **Enter**, the LED changes its state to Action state (LED flashes fast). Then add the adapter to the system. When finished, press **Enter** again to turn on 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).

Removal of adapters requires deconfiguration in AIX first. The adapter must be in defined state or removed from the ODM.

Figure 8 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,scsi3
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 8. Example - Adding a Hot-Plug PCI Adapter*

### Remote I/O Ports

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 6H1 systems with non-looped configurations could experience degraded performance and serviceability. If a non-loop connection is detected, a problem is reported.

The Model 6H1 CEC contains two RIO ports capable of supporting up to two I/O drawers. A one I/O drawer Model 6H1 system cables the primary I/O drawer to both RIO ports. A two drawer system cables one drawer to each RIO port of the CEC with an additional cable connecting the two I/O drawers.

The Model 6H1 ports are designated RIO-0 and RIO-1. Each port operates in bi-directional mode and is capable of passing up to eight bits of data each direction on each cycle of the RIO port. On the Model 6H1, the ports operate at 500 MHz. Based on this, the maximum data rate is 2 GB/s with both RIO ports.

The primary I/O drawer must be attached to the RIO-0 port of the CEC for system bring-up and system support processor access. A single drawer Model 6H1 system will utilize RIO-0 (500 MB/s bi-directional, in total 1 GB/s). A dual drawer Model 6H1 system will utilize both RIO ports (2 x 500 MB/s bi-directional, in total 2 GB/s)

### Considerations for a Second I/O Drawer

A secondary I/O drawer provides an additional fourteen hot-plug capable PCI slots and two additional media bays (no additional connectors for external devices).

> **Note**
>
> Requires RPQ 8A1140 for systems equipped with 450 MHz RS64 III processor features #5201, #5202, or #5204.

### Cabling

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

- JTAG[4] (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 or three per system)

  Remote I/O cables connect the CEC with the I/O drawer(s). If the system is configured with only one I/O drawer, two RIO cables are required. If another I/O drawer is added to the system, a third RIO cable is required to build the loop. RIO cables can be ordered in three lengths (3 m # 3142, 6 m # 3143, 15 m # 3144).

When adding a secondary I/O drawer, a fourth type of interface cable is provided:

- SPCN[5] Cables (two per system)

  SPCN cables are used for power control of the secondary I/O drawer. These cables connect the primary with the secondary I/O drawer and can also be ordered in 3 lengths (3 m # 6006, 6 m # 6008, 15 m # 6007). These cables are the same as those used in Model S80.

Figure 9 on page 17 shows the cabling of a one and two I/O drawer system

[4] Joint Test Action Group (JTAG)
[5] Serial Power Control Network (SPCN)

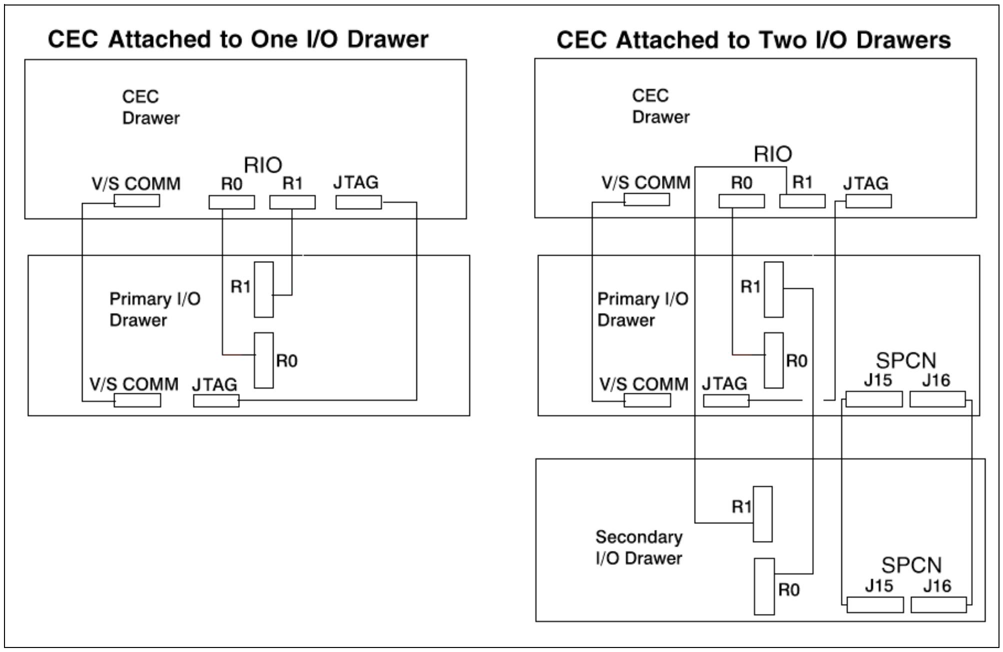


*Figure 9. Cabling of the Drawers*

**Dual Boot Bay Option**

For booting from disks, external disk storage is required for the Model 6H1 as 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-pluggable) 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 prevents 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 disks can be connected to an external SCSI adapter port using the external SCSI adapter to internal IPL disk bay cable assembly (# 3139) and appropriate external SCSI cabling. The external SCSI adapter to internal IPL disk bay 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 converter cable, VHDCI to P, mini 68-pin to 68-pin, 0.3 m (# 2118) 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 10 on page 18 shows how the dual boot bay fits into the I/O drawer.

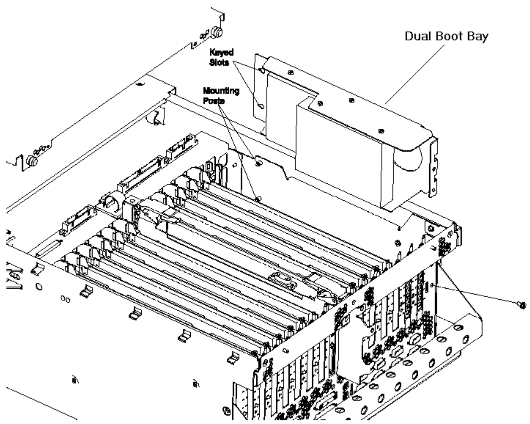


*Figure 10. Dual Disk Bay Option*

### Software Requirements

The Model 6H1 requires AIX Version 4.3.3 with the AIX 4330-10 recommended maintenance package (APAR IY17356), which is included on all pre-installed systems and on the 04/2002 Update CD that ships with AIX Version 4.3.3 as of April 2002 or AIX 5L Version 5.1 with 5100-02 maintenance package or later.

You can also download the actual maintenance level from the Internet to install the machine using NIM[6]. The URL to obtain the maintenance level is:

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

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

## 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 RS/6000 servers. The HACMP solution exploits redundancy between server resources and provides application uptime. The Model 6H1 is available in a high-availability cluster solution package named the HA-6H1. This solution consists of the following components:

- Two Model 7026-6H1 Enterprise 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)
- One 7133-D40 SSA disk subsystem with at least four disk drives
- One 7014-T00 system rack
- All necessary redundant hardware and cables

[6] Network Installation Manager (NIM)

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

## Reliability, Availability, and Serviceability (RAS) Features

Some RAS features, such as redundant power supplies or N+1 hot-plug fans, have been discussed already. Additional topics are covered in the following sections.

### Error Recovery for Caches and Memory

The RS64 IV processor L1 cache, the L2 cache, system busses, 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 and when memory is configured in quads, 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 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 that performs continuous background reads of data from memory to check for correctable errors. Correctable errors are corrected and rewritten to memory, and a threshold counter is maintained that will signal the service processor with a special attention when the threshold is exceeded.

### Chipkill

The Model 6H1 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 6H1 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 transparent to the system.

**Note**

The Chipkill Memory technology is not supported on the 256 MB (2 X 128 MB) Memory (#4110)

### 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 processor 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 uni-processor systems, would change their mode to uni-processor 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[7] records on the FRU[8] so that information stays with the FRU and gets 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[9]/POST[10] testing 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 processor or memory module (as determined by the service processor)
- It reaches a threshold of recovered failures (for example, ECC correctable L2 cache errors, see the preceding) that result in a predictive call-out (as determined by service processor)

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

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

this frozen mode, DMAs[11] 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 attempt recovery. 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 $I^2C$[12] 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 a hardware event. There are different warnings, such as cooling warnings or power fail warnings, that 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 shut down.
- 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 6H1 has an integrated 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 pSeries and RS/6000. This means that all service processor menu functions (using 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 discuss selected features of the 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.

[11] Direct Memory Access (DMA)
[12] Inter-IC ($I^2C$)

- 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 preprogrammed 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 (if, for example, 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 6H1 Service Guide*, SA38-0566.

> **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 the default, allows selection of the IPL type, mode, and speed for your boot capabilities using service processor menus. When selecting fast boot, this 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 Model 6H1 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. As 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 you to select, among other things, to boot to 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.

## System Upgrades

For owners of a 7026-H80, it is possible to upgrade to equivalent processing power of a Model 6H1 while keeping the original serial number.

Model H80 system's processors can be converted to Model 6H1 processors. Table 2 provides the different feature conversions.

*Table 2. Conversion matrix*

| Existing configuration | Required configuration |
|---|---|
| RS64 III, 1-way, 450 MHz | RS64 III, 2-way, 450 MHz |
| RS64 III, 1-way, 450 MHz | RS64 III, 4-way, 450 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 2-way, 600 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 4-way, 600 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 6-way, 668 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 2-way, 750 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 III, 1-way, 450 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 III, 2-way, 450 MHz | RS64 III, 4-way, 450 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 2-way, 600 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 4-way, 600 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 6-way, 668 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 2-way, 750 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 III, 2-way, 450 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 III, 4-way, 450 MHz | RS64 IV, 4-way, 600 MHz |
| RS64 III, 4-way, 450 MHz | RS64 IV, 6-way, 668 MHz |
| RS64 III, 4-way, 450 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 III, 4-way, 450 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 III, 6-way, 500 MHz | RS64 IV, 6-way, 668 MHz |
| RS64 III, 6-way, 500 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 III, 6-way, 500 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 1-way, 600 MHz | RS64 IV, 2-way, 600 MHz |

[1] Conversions are not available for systems equipped with secondary I/O drawer features 6325 and 6327

| Existing configuration | Required configuration |
|---|---|
| RS64 IV, 1-way, 600 MHz | RS64 IV, 4-way, 600 MHz |
| RS64 IV, 1-way, 600 MHz | RS64 IV, 6-way, 688 MHz |
| RS64 IV, 1-way, 600 MHz | RS64 IV, 2-way, 750 MHz |
| RS64 IV, 1-way, 600 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 IV, 1-way, 600 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 2-way, 600 MHz | RS64 IV, 4-way, 600 MHz |
| RS64 IV, 2-way, 600 MHz | RS64 IV, 6-way, 688 MHz |
| RS64 IV, 2-way, 600 MHz | RS64 IV, 2-way, 750 MHz |
| RS64 IV, 2-way, 600 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 IV, 2-way, 600 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 4-way, 600 MHz | RS64 IV, 6-way, 688 MHz |
| RS64 IV, 4-way, 600 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 IV, 4-way, 600 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 6-way, 688 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 1-way, 750 MHz | RS64 IV, 2-way, 750 MHz |
| RS64 IV, 1-way, 750 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 IV, 1-way, 750 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 2-way, 750 MHz | RS64 IV, 4-way, 750 MHz |
| RS64 IV, 2-way, 750 MHz | RS64 IV, 6-way, 750 MHz |
| RS64 IV, 4-way, 750 MHz | RS64 IV, 6-way, 750 MHz |

[1] Conversions are not available for systems equipped with secondary I/O drawer features 6325 and 6327

The existing Model H80 processor being replaced should be returned to IBM.

Memory DIMMs from the H80, adapters, and most of the disk drives can remain compatible. For further information about adapters, especially if the adapter is supported in Model 6H1, refer to the *PCI Adapter Placement Reference Guide*, SA38-0538.

When upgrading from an Model H80 it is required to upgrade to AIX 4.3.3 ML10 or higher before you can run the 6H1 system. When migrating from Model H70 in a 7015-R00, a new rack is needed because the 7015-R00 is not a supported configuration for a Model 6H1. An existing 7014-S00 rack can be used when installing a Model 6H1, except if the S00 rack is used with DC power.

## External Storage Expandability

The storage requirement for the Model 6H1 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 6H1 can be attached to IBMs Enterprise Storage Server. And, by using the Fibre Channel Adapter, the Model 6H1 can be attached to the 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 H80, 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 2 SCSI-3 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 contains one or two DLT tape drives providing 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, which includes 3590-E11 (two requires 12 EIA), 3590-B11 (two requires 12 EIA), 3570-C12 (6 EIA), 3570-C11 (6 EIA), and 3490-F11 (4 EIA).

### SP Attachment and Clustering

The joining of the Model 6H1 to the RS/6000 SP satisfies the need many SP environments have for large, powerful, and memory rich processors for their database servers and today’s ebusiness 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 that is essential to resolving 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 thirty two Model 6H1 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 6H1 servers centrally managed through Parallel Systems 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 6H1 Installation Guide,* SA38-0575
- *IBM @server pSeries 660 Model 6H1 User's Guide*, SA38-0565
- *IBM @server pSeries 660 Model 6H1 Service Guide*, SA38-0566
- *PCI Adapter Placement Reference Guide*, SA38-0538
- 7014 Series Model T00 and T42 Rack Installation and Service Guide, SA38-0577
- Site Hardware Planning Information, SA38-0508

## Select IBM Redbooks

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

- *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/aix/
www.chips.ibm.com/
www.research.ibm.com/topics/serious/chip/
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!

| | |
|---|---|
| Amy, Larry | IBM Austin |
| Beebe, Bill | IBM Austin |
| Enders, Rich | IBM Austin |
| Frey, Brad | IBM Austin |
| Haug, Volker | IBM Germany |
| Henderson, Daniel | IBM Austin |
| Keenan, Ernest A. | IBM Austin |
| Lee, Dennis | IBM Austin |

| | |
|---|---|
| Lutz, Stephen | IBM Germany |
| Manohar, Shyam | IBM India |
| McCord, Scott | IBM Austin |
| Ruth, Dave | IBM Austin |
| Scherrer, Carolyn | IBM Austin |
| Smylie, Scott | IBM Austin |
| Shempert, Craig | IBM Austin |
| Stys, Mike | IBM Austin |
| Toutant, Ed | IBM Austin |
| Venton, Todd | IBM Austin |

## Biography

**Gregor Linzmeier** is an IT Specialist in the technical support for RS/6000 and pSeries as part of the Web Server Sales Organization in Mainz, Germany. Gregor is a member of the Workstation Technology Focus Group, supporting IBM sales, Business Partners, and customers with pre-sales consultation and implementation of client/server environments. He has worked for more than thirteen years as an infrastructure specialist for RT, RS/6000, and AIX in large CATIA client/server projects.

## Special Notices

This document was produced in the United States. IBM may not offer the products, programs, services, or features discussed herein in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products, programs, services, and features available in your area. Any reference to an IBM product, program, service, or feature is not intended to state or imply that only IBMs product, program, service, or feature may be used. Any functionally equivalent product, program, service, or feature that does not infringe on any of IBMs intellectual property rights may be used instead of the IBM product, program, service, or feature.

Information in this document concerning non-IBM products was obtained from the suppliers of these products, published announcement material, or other publicly available sources. Sources for non-IBM list prices and performance numbers are taken from publicly available information including D.H. Brown, vendor announcements, vendor WWW Home Pages, SPEC Home Page, GPC (Graphics Processing Council) Home Page, and TPC (Transaction Processing Performance Council) Home Page. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBMs future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of a specific Statement of General Direction.

The information contained in this document has not been submitted to any formal IBM test and is distributed "AS IS". 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. The use of this information or the implementation of any techniques described herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.

IBM is not responsible for printing errors in this publication that result in pricing or information inaccuracies.

The information contained in this document represents the current views of IBM on the issues discussed as of the date of publication. IBM cannot guarantee the accuracy of any information presented after the date of publication.

All prices shown are IBMs suggested list prices; dealer prices may vary.

IBM products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information provided in this document and information contained on IBMs past and present Year 2000 Internet Web site pages regarding products and services offered by IBM and its subsidiaries are "Year 2000 Readiness Disclosures" under the Year 2000 Information and Readiness Disclosure Act of 1998, a U.S statute enacted on October 19, 1998. IBMs Year 2000 Internet Web site pages have been and will continue to be our primary mechanism for communicating year 2000 information. Please see the "legal" icon on IBMs Year 2000 Web site (www.ibm.com/year2000) for further information regarding this statute and its applicability to IBM.

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

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: ADSTAR, AIX, AIX/6000, AIXwindows, AS/400, C Set++, CICS, CICS/6000, DB2, ESCON, IBM, Information Warehouse, Intellistation, LANStreamer, LoadLeveler, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, OS/2, OS/400, OS/390, Parallel Sysplex, POWERparallel, RS/6000, S/390, Service Director, System/390, ThinkPad, TURBOWAYS, VisualAge. The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: AIX PVMe, AS/400e, DB2 OLAP Server, DB2 Universal Database, DEEP BLUE, e-business (logo), eNetwork, eServer, GigaProcessor, HACMP/6000, Intelligent Miner, Network Station, pSeries, POWER2 Architecture, PowerPC 604, SmoothStart, SP, Videocharger, Visualization Data Explorer, WebSphere. A full list of U.S. trademarks owned by IBM may be found at http://iplswww.nas.ibm.com/wpts/trademarks/trademar.htm.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States, other countries, or both. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group. Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Lotus and Lotus Notes are trademarks or registered trademarks of Lotus Development Corporation. Tivoli, TME, TME 10, and TME 10 Global Enterprise Manager are trademarks of Tivoli Systems, Inc. Other company, product, and service names may be trademarks or service marks of others.