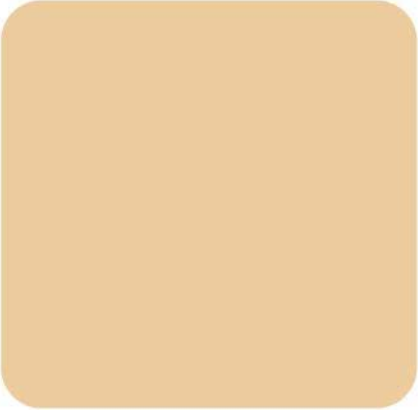




**an hp technical  
white paper**

september 2002



**hp server rp7405**

## table of contents

<b>introduction</b> .....	<b>4</b>
enterprise-class high availability .....	4
industry-leading functionality .....	4
easy ordering and deployment.....	4
<b>right performance, right price: the hp server rp7405</b> .....	<b>5</b>
key features .....	5
unmatched performance and scalability .....	5
superior operating system .....	6
flexible configuration .....	6
<b>hp server rp7405 system overview</b> .....	<b>7</b>
rp7405 front view .....	7
rp7405 rear view .....	8
<b>hp server rp7405 features at a glance</b> .....	<b>9</b>
hp server rp7405 system architecture.....	9
cell boards .....	10
cell design details .....	10
cell configurations.....	10
cell memory .....	11
memory latencies.....	11
cell hot-plug .....	12
superior processor cell cooling with the turbo-cooler .....	12
I/O subsystem.....	13
I/O controller chips .....	13
core I/O cards .....	13
PCI backplane.....	14
supported cards .....	15
core I/O .....	15
management processor .....	17
external LAN port .....	17
external Ultra160 SCSI port.....	17
internal peripheral bay.....	17
hot-plug disk drives .....	17
removable-media bay .....	18
I/O configuration guidelines.....	18
<b>ac power subsystem</b> .....	<b>18</b>
requirements for fault-tolerant power compliance .....	18
fault-tolerant power compliance in the rp7405 .....	19
features of the ac power system.....	20
20-ampere circuits .....	20
dual ac modular bulk power supplies (BPS).....	20
ac power consumption .....	20
maximum and typical power consumption values .....	20
60-ampere power distribution unit .....	21
30-ampere power distribution unit .....	23
16-ampere power distribution unit .....	24
<b>performance and scalability</b> .....	<b>26</b>
I/O bandwidth.....	26
scalability .....	26
performance considerations and configuration rules .....	26
<b>racking</b> .....	<b>27</b>
ballast kits .....	27
third-party racks .....	27

<b>nPartitions .....</b>	<b>28</b>
how partitions work .....	28
nPartition configurations .....	28
reconfiguring partitions .....	28
dynamic reconfiguration .....	28
uses for nPartitions .....	29
more details about some key uses for nPartitions .....	29
resource partitioning .....	29
multiple environments .....	30
improved availability .....	30
more flexible scaling .....	30
what you gain with nPartitions .....	30
<b>high availability.....</b>	<b>31</b>
partition reliability .....	31
CPU protection .....	31
memory protection .....	33
protection for I/O and disks.....	33
reliability in the cabinet infrastructure .....	34
dual ac line cord support.....	34
redundant power supplies.....	34
resilience to service processor failures .....	34
<b>serviceability .....</b>	<b>35</b>
enhanced serviceability .....	35
easy access to all components.....	35
attractive, functional design .....	36
<b>ease of management .....</b>	<b>37</b>
rp7405 management processor .....	37
partition manager software .....	37
servicecontrol suite.....	37
<b>operating environment.....</b>	<b>38</b>
hp-ux 11i support .....	38
key hp-ux 11i features.....	38
choice of hp-ux operating environments.....	39
<b>hp services for the hp server rp7405 .....</b>	<b>39</b>
design, migrate, and integrate .....	39
consulting services .....	39
deployment services .....	39
hp startup integration express service.....	40
hp smartest integration service .....	40
hp onsite partitioning service.....	40
hp educational services .....	40
run and evolve .....	40
standard support and industry-leading mission critical services.....	40
critical systems support.....	40
business continuity support.....	41
support for data storage .....	41
business recovery services .....	41
outsourcing services .....	41
financing .....	41
<b>for more information.....</b>	<b>41</b>

## **introduction**

Today's business environment is changing in a fundamental way. Despite significant competitive pressures, businesses must increasingly stretch their IT budgets further. For many IT managers, this has created a challenging situation: whether to purchase systems that offer enterprise-level availability and functionality—or choose more economical systems that sacrifice functionality and reliability. Regardless of the decision, many businesses find themselves making an uncomfortable trade-off between functionality and price. As an organization's reliance on its IT infrastructure increases, the importance of this trade-off decision can literally be a matter of business survival. Sound familiar?

The HP 05 series servers address this new business reality by offering your business the ability to deliver top-line results without compromising on the bottom line. Designed for workgroup applications; front-end applications (such as caching, firewall, and workload balancing); and Web, application, and small database serving, 05 series systems offer the right performance while delivering the high-availability and manageability features you've come to expect from HP servers.

## **enterprise-class high availability**

The 05 series servers from HP allow you to stretch your budget dollars further while offering enterprise-class high availability that no one else can match. Standard high-availability features on all 05 series servers include dynamic processor resilience; full error-checking and-correcting (ECC) protection; and HP's Event Monitoring Service (EMS), which provides proactive fault avoidance, detection, and notification. Select 05 series models also offer dynamic hardware partitioning for full electrical isolation. And the 05 series product family offers a wide range of high-availability solutions to keep your business running, including industry-leading HP MC/Serviceguard.

## **industry-leading functionality**

The 05 series products offer unparalleled opportunity to manage your server resources. Integrated features, such as HP Secure Web Console, give you full remote management capability, including centralized single-system and multi-system configuration. Other industry-leading management tools include HP-UX Workload Manager (WLM), which allows various workloads to be managed to meet the service-level objectives of your business. And select 05 series models offer dynamic virtual partitioning, which allows you to run multiple workloads on unique instances of HP-UX, ensuring complete software fault isolation.

## **easy ordering and deployment**

Designed for ease of ordering and deployment, the 05 series product line has a full range of standardized configurations with the flexibility to meet immediate business requirements. Preconfigured with high-powered PA-RISC processors and offering a choice of four pre-integrated and tested HP-UX 11i operating environments, 05 series servers can be put to work to rapidly address business demands. 05 series systems also offer industry-leading form factors that allow multiple servers to be racked in the data center, preserving valuable floor space and lowering operating costs.

**right performance,  
right price: the  
hp server rp7405**

**figure 1. Front view of the HP Server rp7405.**



### **key features**

The HP Server rp7405 supports from two to eight PA-8700 650 MHz processors. The PA-8700 is one of the latest in a long line of award-winning PA-RISC processors developed by HP.

The rp7405 supports 15 hot-plug PCI cards and up to 32 gigabytes of main memory. It's packed with high-availability features—like HP's midrange and high-end Superdome servers, the rp7405 supports hardware and software partitioning, plus online partition maintenance and upgrades (in dual-partitioned systems). It also has redundant and hot-swappable cooling units and power supplies, hot-plug disks and I/O cards, and an extremely reliable and fault-resilient main memory and cache memory subsystem.

### **unmatched performance and scalability**

The performance and scalability of the rp7405 are unmatched in this class of servers:

- High-performance RISC processor
  - 650 MHz (PA-8700)
  - 4-way superscalar
  - 56-way instruction reorder buffer
  - 10 functional units
  - 2.25 MB on-chip cache (1.5 MB data cache, 0.75 MB instruction cache)
- Outstanding memory subsystem performance
  - Peak memory bandwidth to 8 GB/s
  - High-density system packaging provides 32 GB of total memory capacity (64 GB in the near future)
  - Configurations available from 2 to 8 CPUs
- Superior I/O interface performance; more I/O connectivity
  - 64-bit 66 MHz (4X) industry-standard PCI
  - 15 independent PCI card slots
  - Independent PCI buses for better bandwidth and error containment
  - 8.5 GB/s aggregate I/O slot bandwidth

## superior operating system

Like existing HP UNIX® servers, the rp7405 runs the industry-leading HP-UX 11i operating environment. HP-UX addresses the major computing challenges that customers face today in online transaction processing (OLTP), enterprise resource planning (ERP), supply chain management (SCM), server consolidation, telco billing applications, high-performance technical computing, and customer relationship management (CRM). HP-UX is also ideal for use in business intelligence systems and in Internet, scientific, and technical applications. This mainframe-class, 64-bit operating system enjoys the industry's greatest support from independent software vendors, allowing a choice of more than 15,000 applications—including native 64-bit versions of all major databases and leading ERP applications.

## flexible configuration

Three rp7405 standard configurations are available to choose from. Whether a two-CPU entry-level, four-CPU average-workload, or eight-CPU high-performance configuration is required, there is a preconfigured rp7405 solution ready for the challenge.

rp7405 standard configurations	A7111A	A7112A	A7113A
Number of processors (included)	2	4	8
Maximum processor capacity	8	8	8
Processor	PA-8700 650 MHz	PA-8700 650 MHz	PA-8700 650 MHz
Cell boards (included)	1	2	2
Memory (included)	4 GB	8 GB	16 GB
Max memory capacity	32 GB	32 GB	32 GB
Available/Max I/O card slots	7/15	14/14	14/14
Internal DVD (included)	1	1	1
Internal disks (included, max 4)	2 x 73 GB	4 x 73 GB	4 x 73 GB
Maximum internal disk capacity	292 GB	292 GB	292 GB
Supported OS version	HP-UX 11i	HP-UX 11i	HP-UX 11i


## hp server rp7405 system overview

### rp7405 front view

The following illustrations show major components of the HP Server rp7405 as well as the system's mechanical and architectural features.

**Figure 2** shows a front view of the rp7405 with its front plastic bezel and top and left side panels removed.

**figure 2. Front and side view of the rp7405.**




A peripheral bay located at the top front of the rp7405 provides space for four hot-plug disk drives and one hot-plug removable-media device. Directly below the peripheral bay are two PCI power bricks, which supply dc power for the PCI backplane. Below the power supplies are two redundant hot-swappable cooling fans. These fans pull cool air in from the front and force air rearward, cooling the system's internal components. At the bottom is the bulk power supply bay, which houses two redundant (2N) hot-swap power supplies with dual grid support.

The illustration also shows the right side of the rp7405, with a view of the cell board bay. This bay supports up to two cell boards; the cell boards contain processors, memory, and cell controller chips.

## rp7405 rear view

**Figure 3** is a rear view of the rp7405 showing the location of the two hot-swappable 150 mm exhaust fans and the I/O bay bulkhead directly above them. The core I/O cards are located at the right edge of the unit.

**figure 3. Rear view of the rp7405.**



The system backplane board houses the linkages used for communications between cell boards, I/O, and internal peripherals. The rp7405's I/O card bay is located at the top rear of the system. It contains 15 PCI card slots, all supporting hot-plug functionality.

The bottom rear of the rp7405 has inputs for the 2+2 redundant line cords. Because of the high degree of connectivity offered in the rp7405, a cable management arm (not shown) is provided for dressing cables and simplifying cable routing.

## hp server rp7405 features at a glance

### features and capacities

- 1–2 CPU/memory cell boards, hot-plug
- 2–8 PA-RISC 64-bit CPUs (PA-8700 @ 650 MHz)
- Up to 32 GB of memory
- 15 PCI card slots (66 MHz x 64-bit)
- 2 core I/O cards
- Built-in management processor
- Integrated Web console
- Up to 4 internal hot-plug disk drives (18, 36, and 73 GB)
- 1 hot-plug removable-media bay (holds standard DVD or optional DAT drive)

### speeds and feeds

- 530 MB/s PCI link bandwidth per slot (14 of 15 slots)
- 8.5 GB/s I/O slot bandwidth (peak)
- 8 GB/s bandwidth (peak) for cell-controller-to-memory-subsystem
- 8 GB/s cell to cell link (peak)

### partitioning

- 1–2 nPartitions in hardware
- Up to 8 virtual partitions (vPars)<sup>1</sup>

### high availability

- 2N redundant hot-swap power supplies
- Redundant hot-swap fans, all with HP Smartfan technology
- Redundant power line inputs for dual grid coverage
- Error checking and correcting (ECC) on all CPU and memory paths
- ECC on all system cache memory
- Main memory DRAM kill resiliency (“chip kill”)
- Redundant dc-dc converters for key subsystems
- Parity-protected I/O data paths
- Independent I/O paths

### operating system

- HP-UX 11i operating environment (IPR 0203 or later)

### physical characteristics

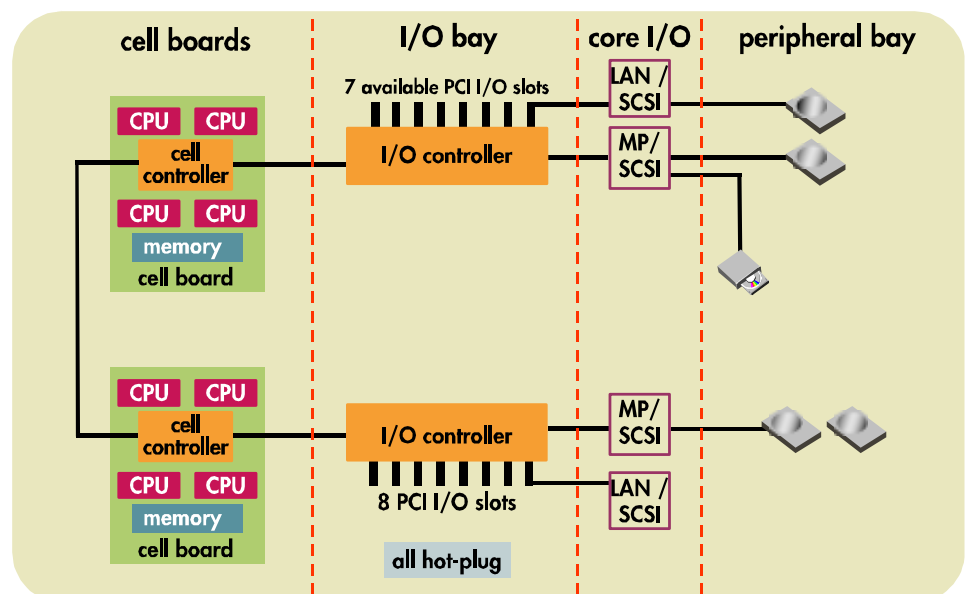
- Rack-mount configuration only
- Height: 10 EIA units (17.5 in. [44.45 cm])
- Depth: 30 in. (76.20 cm)
- Width: 19 in. (48.26 cm)

<sup>1</sup>Virtual partitions supported late 2002

## hp server rp7405 system architecture

The HP Server rp7405 architecture can be thought of as a set of modular “building blocks” that fit together to form a high-performance, scalable, highly available and flexible computing platform. The rp7405 is designed around a cell-based architecture that allows it to be configured as one symmetrical multi-processing (SMP) server or as two hardware-independent partitions.


figure 4. Architecture of the rp7405, showing the basic modular “building blocks” of the system.



## cell boards

The cell board is one of the three basic building blocks of the HP Server rp7405. A cell, or cell board, is a hot-plug module that primarily houses processors and memory. **Figure 5** shows the layout of the rp7405 PA-8700 cell board.

**figure 5. Layout of an rp7405 PA-8700 cell board.**



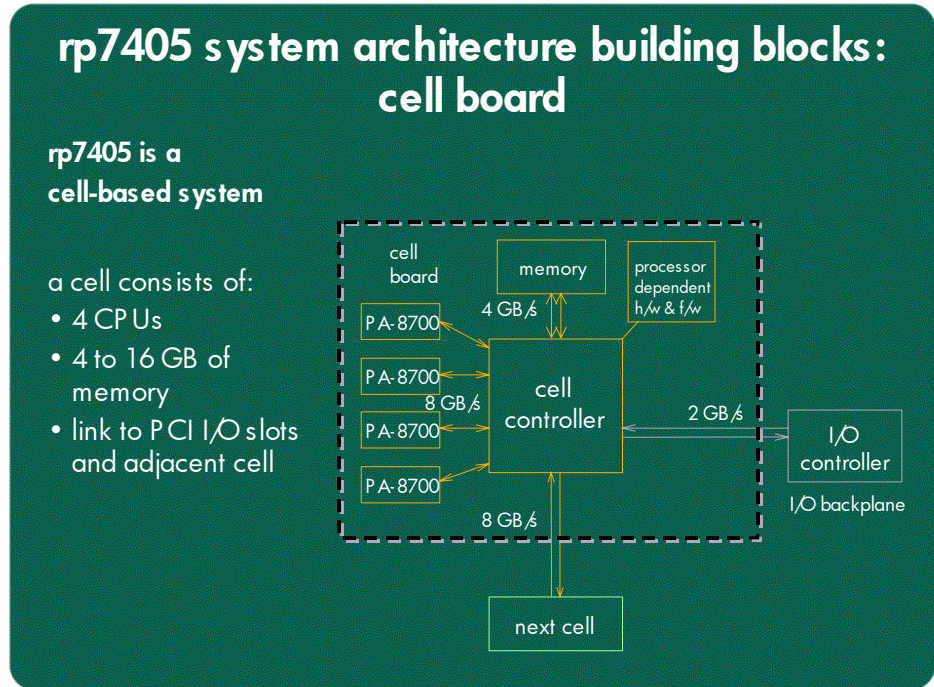
## cell design details

The cell controller (CC) is at the heart of each cell board. The CC provides the communications link between processors, memory, I/O, PDH (processor-dependent hardware), and the adjacent cell. Surrounding the CC are up to four processors and up to 16 GB of main memory. Cell-to-cell communication is at 8 GB/s of peak bandwidth, and CC-to-I/O-subsystem bandwidth is 2 GB/s peak. The cell controller chip contains interface logic, and it also maintains cache coherency throughout the system.

## cell configurations

An rp7405 can contain a minimum of one and a maximum of two cells. Each cell can be configured with two or four active processors.

figure 6. An rp7405 cell supports up to four processors and up to 16 GB of memory.



#### cell memory

The rp7405 standard configurations include pre-integrated memory. The minimum cell configuration includes two active CPUs and 4 GB of memory per cell board. The maximum rp7405 standard configuration includes four active CPUs and 8 GB of memory per cell board. Additional memory DIMM modules for the rp7405 are sold in sets of four (quads), with DIMM board sizes of 256 MB, 512 MB, and 1 GB available. The maximum memory configuration is 16 GB per cell board. Memory quads of different sizes can be mixed within a chassis and within a cell. However, for optimum memory interleaving and performance, it is recommended that one memory size be selected and distributed evenly across available cells.

Within a cell the CC-to-memory peak bandwidth is 4 GB/s. Memory is accessed directly through the CC; therefore, all memory slots are accessed regardless of the number of processors in the cell.

#### memory latencies

There are two types of memory latencies within the rp7405 system:

- **memory latency within the cell**—This latency refers to the case where an application runs on a partition that consists of a single cell.
- **memory latency between cells**—This type of latency is present when the partition consists of two cells. In this case half of the addresses are to the memory on the requesting processor's cell, and the other half of the addresses are to the memory of the other cell.

Memory latency in the rp7405 depends on the number of CPUs and the location of their corresponding cell boards. Assuming that there is equally distributed access to all memory controllers and that cell boards are installed to minimize latency, the average load-to-use latency on an idle rp7405 is shown in this table:

number of CPUs	average idle memory latency
4 CPUs	220 ns
8 CPUs	250 ns

## cell hot-plug


At its initial release, the HP Server rp7405 supports cell hot-plug. Coupled with the server's partitioning capability\*, cell hot-plug allows for the servicing of a cell board within a single partition while the other partition continues normal operation. Any number of configuration changes can be made to the partition being serviced, including replacing the complete cell board and adding or deleting CPUs and memory. This means that a cell can be added "on-line" (to a single-cell rp7405) to increase the computing resources of the rp7405. (Cell hot-plug is supported only in systems configured with two partitions.)

## superior processor cell cooling with the turbo-cooler

As one of the world's most powerful microprocessors, the PA-RISC processor and its associated core electronics require an efficient mechanism to remove excess heat. Removing that heat is important to one of the primary design objectives for the rp7405—delivering industry-leading performance density while maintaining high levels of availability and reliability. Efficient heat removal is one of the key techniques HP uses to ensure the highest levels of reliability in all of the components that make up our server computers.

The key to processor cell cooling in the rp7405 is the turbo-cooler. The turbo-cooler allows new levels of performance density and efficiency while maintaining the high levels of reliability and availability that HP customers have come to expect. The improved density provided by turbo-coolers is apparent in the compact packaging of the rp7405 cell assembly.

**figure 7. The turbo-cooler provides quiet, efficient, and highly reliable cooling.**

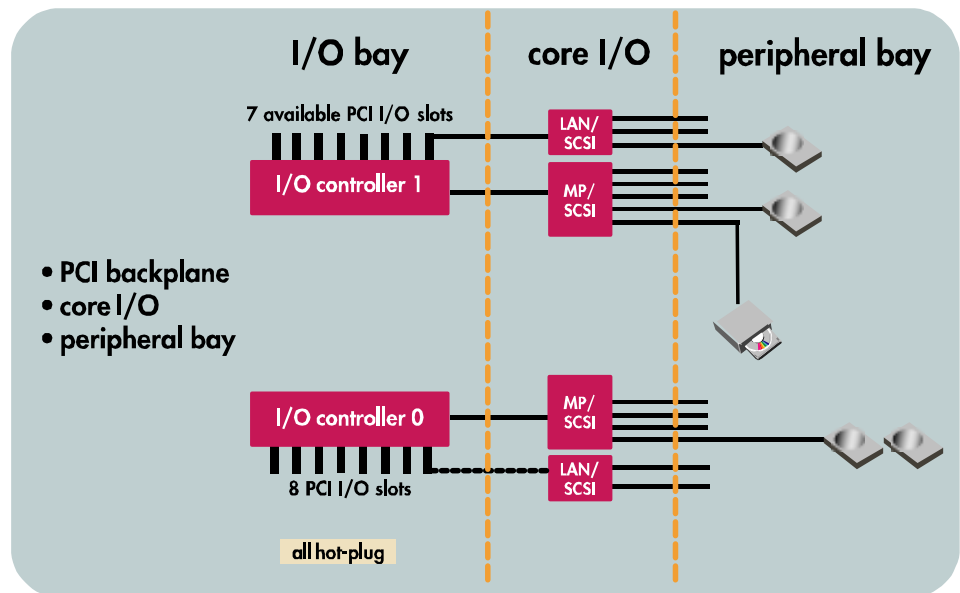


\* The rp7405 can be configured as a single large SMP server or hardware-partitioned into two smaller logical servers. See "nPartitions" on page 29 for more details about partitioning.

## I/O subsystem

The I/O subsystem is one of the building blocks of the rp7405. Components within the I/O subsystem are the I/O controllers, internal peripheral bay, and multifunction core I/O. **Figure 8** shows a basic diagram of the I/O subsystem.

**figure 8.** Basic block diagram of the rp7405 I/O subsystem.



## I/O controller chips

The rp7405 contains two master I/O controller chips located on the PCI backplane. Each I/O controller contains 16 high-performance, 12-bit-wide links; these links connect to 18 slave I/O controller chips supporting the PCI card slots and core I/O.

In the rp7405, two links—one from each master controller—are routed through the system backplane and are dedicated to core I/O. The remaining 30 links are divided among the sixteen “4X” 66 MHz x 64-bit PCI card slots\*, with each slot on a PCI bus by itself. This one-card-per-bus architecture leads to greater I/O performance, better error containment, and higher availability.

Each controller chip is also directly linked to a host cell board. This means that both cell boards must be purchased in order to access all 15 available I/O card slots. (With only one cell board, the system has access to seven expansion slots.)

## core I/O cards

The rp7405 can be purchased with either one or two core I/O card sets. Each core I/O product contains two cards: an MP/SCSI card and a LAN/SCSI card. The core I/O cards provide console, Ultra160 SCSI, Gigabit LAN, serial, and management processor functionality. The second core I/O card set can be used to enable dual partitioning in the rp7405 and provide access to a second set of disk drives.

The LAN/SCSI card provided with each core I/O product occupies one of the 16 PCI slots. Since there must always be at least one core I/O card set, the rp7405 has 15 PCI slots available for expansion cards. If the second core I/O product is included (A7112A/A7113A), there are fourteen remaining slots available for expansion cards.


The internal peripheral bay supports up to four low-profile disks and one removable-media device. The internal disks are electrically divided into two pairs. SCSI controller chips located on each core I/O card set support each pair of internal disks. This means that both core I/O card sets must be present in order to access both halves of the peripheral bay. This also means that I/O paths are not shared and disks are electrically isolated, providing for optimal isolation between partitions.

\* 66 MHz x 64-bit slots are also known as “4X” slots. Each core I/O consumes one PCI slot. Therefore, if two core I/O card sets are included (A7112A and A7113A), then there are 14 slots available.

**Figure 9** shows a more detailed view of the PCI backplane. Note that 14 of the 15 available expansion slots are supported by dual high-performance links. These dual-link I/O slots are called “twin-turbo” slots because both links operate at turbo speeds—that is, each offers 265 MB/s of bandwidth—providing a maximum of 530 MB/s (peak) bandwidth for the slot. The remaining two I/O slots are single links and are designated as “turbo” slots. Aggregate I/O slot bandwidth is 8.5 GB/s.

Every PCI slot in the rp7405 is capable of running at 66 MHz x 64 bits—known as “4X” PCI by HP. This means that each rp7405 I/O slot will allow the industry’s highest-performing PCI cards to run at their maximum design speed.

**figure 9. The rp7405’s PCI backplane has dual high-performance links for 14 of the 16 I/O card slots.**



In practice, PCI I/O cards requiring the most bandwidth should be configured into the twin turbo slots. Because each I/O slot has a dedicated bus, the card in any slot can be hot-plugged or serviced without affecting other slots.

The rp7405 I/O card slots are keyed for either 3.3-volt or 5-volt PCI cards. There are 12 slots keyed for 3.3 volts, and these support either 3.3-volt or universal PCI cards. The remaining four slots are keyed for 5 volts and support either 5-volt or universal PCI cards.

All cards on the rp7405’s list of supported cards are universally keyed. Therefore, any supported card will work in any slot.

## supported cards

HP strives to provide the latest in I/O technology, offering customers greater performance and connectivity—and the rp7405 supports an extensive list of industry-standard PCI I/O cards. The list below shows the list of I/O cards supported at the first release of the rp7405. The maximum number of each card per system is 15, except as noted.

### mass storage

- 2 GB Fibre Channel
- Single-port PCI 2X Fibre Channel
- Single-port Ultra2 SCSI HBA
- Ultra2 SCSI RAID (maximum 12 cards per system)
- Dual-port F/W SCSI-2
- Dual-port Ultra2 SCSI
- PCI FWD SCSI-2
- PCI Ultra160 SCSI
- PCI dual-channel Ultra160 SCSI

### local area network (LAN)

- Single-port 1000Base-SX Ethernet LAN
- Single-port 1000Base-T Ethernet
- 4-port 100Base-TX Ethernet
- Single-port 100Base-TX Ethernet
- FDDI dual-attach station
- ATM 622 Mbps MMF
- ATM 155 Mbps UTP or MMF
- 802.5 Token Ring 4, 16, 100 Mbps operation

### combination cards

- 2-port Ultra2 SCSI/2-port 100Base-T combo card

### wide area network (WAN)

- 2-port X.25/Frame relay
- 4-port X.25/Frame relay

### serial multiplexer cards

- 64-port MUX
- 8-port MUX

### cluster interconnect

- Hyperfabric 4X interconnect (maximum 4 cards per system)
- Hyperfabric 2X interconnect (maximum 4 cards per system)

### encryption

- HP Speedcard

New I/O cards are added regularly. For the latest list of supported I/O cards, please refer to the online version of the *HP Server Configuration Guide*.

## core I/O

The rp7405 chassis supports up to two core I/O card sets. Each set contains two cards (MP/SCSI and LAN/SCSI), which are installed in different locations. MP/SCSI cards are installed along the right rear vertical edge of the chassis; LAN/SCSI cards are installed in the PCI card bay. A minimum of one core I/O card set must be installed in each system; the second core I/O card set can be used to enable hardware partitioning or to utilize the full capacity of the built-in mass storage bays. One core I/O card set is included in A7111A; two core I/O card sets are included in A7112A and A7113A.

Both core I/O card sets are identical. However, the electrical connections to internal peripherals and the I/O controller are slightly different. In the “primary” core I/O card set, the LAN/SCSI board is supported by a single 265 MB/s link; in the “secondary” core I/O card set, the LAN/SCSI board is supported by two 265 MB/s links. In addition, in the primary core I/O set, the two SCSI controllers—one in the LAN/SCSI card and one in the MP/SCSI card—each support a single internal disk drive. In the secondary core I/O, only the MP/SCSI board is used to support disk drives; however, both disk drives are supported off this single SCSI controller and bus.

The performance of the primary LAN/SCSI card will be less than that achieved by the secondary LAN/SCSI. The primary LAN/SCSI card, located in I/O chassis 1 (slot 8), will achieve LAN performance levels of approximately 550 Mb/s (transmit) and 940 Mb/s (receive). The secondary LAN/SCSI card, located in I/O chassis 0, will achieve nominal Gigabit LAN performance. The primary LAN/SCSI card will achieve SCSI performance levels of approximately 65 MB/s outbound (writes to disk) and 120 MB/s inbound (reads from disk). The secondary LAN/SCSI card will achieve nominal Ultra160 SCSI performance.

Figure 10 and figure 11 show the basic block diagrams of primary and secondary core I/O and illustrate the differences.

figure 10. Block diagram showing basic core I/O card set (primary) in the rp7405.

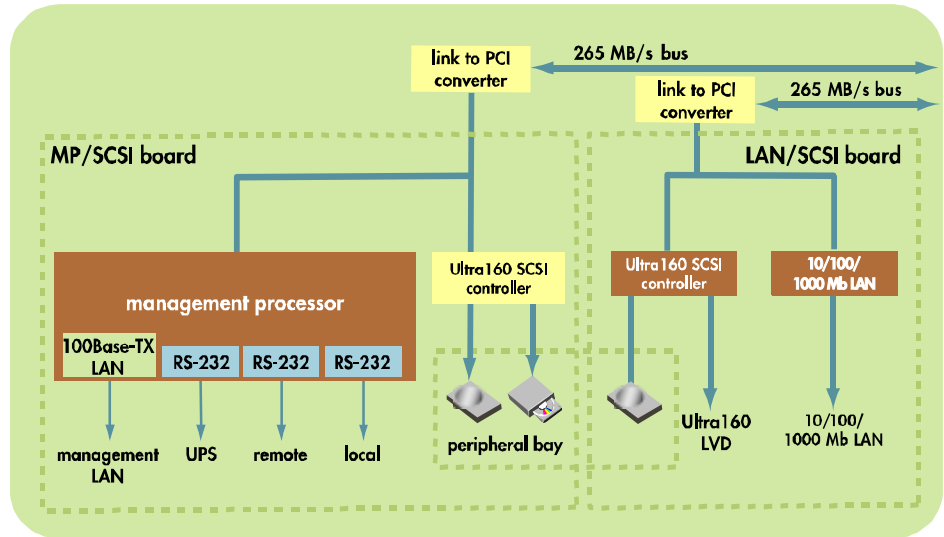
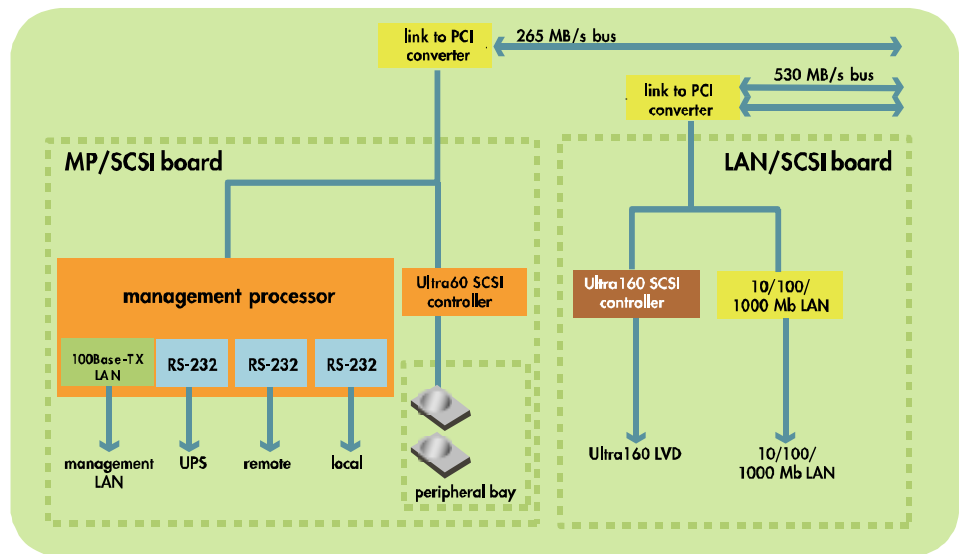


figure 11. Basic core I/O card set (secondary).



management processor	<p>The management processor (MP) is a dedicated processor located on each MP/SCSI card that simplifies and extends system management and also enhances serviceability. The MP minimizes or eliminates the need for the system administrator to be physically at the system to perform tasks such as diagnostics, system management, or even hard resets. Here are some of the features enabled by the rp7405 management processor:</p> <ul style="list-style-type: none"> <li>• System management over the Internet or intranet</li> <li>• System console redirection</li> <li>• Console mirroring</li> <li>• System configuration for automatic restart</li> <li>• Viewing history log of system events</li> <li>• Viewing history log of console activity</li> <li>• Setting MP inactivity timeout thresholds</li> <li>• Remote system control</li> <li>• Remote power cycle (except for MP housekeeping power)</li> <li>• Viewing system status</li> <li>• Event notification to system console, e-mail, pager, and/or HP Response Centers (e-mail and pager notification work in conjunction with HP's Event Monitoring System [EMS])</li> <li>• Automatic hardware protection from critical environmental problems</li> <li>• Access to management interface and consoles on WAN failure (modem required)</li> <li>• Automatic system restart</li> <li>• Remote resetting of hardware partitions</li> <li>• Forward progress indicator (via a virtual front panel)</li> <li>• Out-of-band manageability and processor dependent code (PDC) firmware update</li> <li>• Configuration of manageability and console security</li> <li>• Secure Sockets Layer (SSL)</li> </ul>
external LAN port	The external LAN port is a 10/100/1000Base-T external LAN port using an RJ45 connector.
external Ultra160 SCSI port	The external SCSI port is an Ultra160 LVD external SCSI port for connections to mass storage or media.
<b>internal peripheral bay</b>	<p>The rp7405 internal peripheral bay is located at the top front of the system chassis. The peripheral bay holds up to four low-profile hot-plug disks and one hot-plug removable-media device.</p> <p>Each rp7405 core I/O card set contains dual-channel Ultra160 SCSI controller chips that support the SCSI devices in the internal peripheral bay. Each core I/O card set supports two internal disks. It is important to note that the two disks supported by the primary core I/O card set are managed by separate controllers and SCSI buses. A single controller and SCSI bus manage the second pair of disks supported by the secondary core I/O card set. If use of more than two internal disks is needed, the rp7405 will require both core I/O card sets.</p>
hot-plug disk drives	The rp7405 holds four embedded SCSI disks, which are accessible from the front of the server. These are hot-plug disks, so they can be removed and inserted while the rp7405 continues to operate. The system supports the following disks: 18 GB, 10k rpm; 36 GB, 15k rpm; and 73 GB, 10k rpm. Two or four 73 GB disks are included with each rp7405 standard configuration.

## removable-media bay

The rp7405 contains one removable-media bay, which supports either a DVD-ROM or a DDS-4 DAT drive. All rp7405 products come standard with an integrated DVD-ROM. The DDS-4 DAT drive can be purchased to replace the DVD-ROM (field replacement only). Access to these devices is also from the front. The DVD-ROM drive provides enhanced features while preserving backward read compatibility with CD-ROM drives. Data transfer rates of up to 6.75 MB/s are achieved with the DVD format; 4.8 MB/s can be achieved using the CD-R format. The DDS-4 drive has a maximum storage capacity of 40 GB, with a peak transfer rate of 21.6 GB/hour (compressed).

A dedicated SCSI bus from core I/O supports the removable-media bay. This configuration allows for either the DVD or DAT drive to be hot-plugged without interrupting server operation.

## I/O configuration guidelines

The following table summarizes the configuration rules for the usage of I/O slots and internal peripherals in the rp7405.

configuration	minimum requirements	
	minimum number of cells	minimum number of core I/O sets
>7 I/O card slots or access to both I/O card bays	2	1
>2 internal disks or access to both pairs of disks	2	2
2 partitions	2	2

## ac power subsystem

The HP Server rp7405 contains the exact same AC power subsystem found in the rp7410. The rp7410 was the first-ever server to meet the Uptime Institute's new Fault Tolerant Power Compliance Specification, Version 2.0, with no deviations. And it's only the second midrange server to be certified by The Uptime Institute under *any* Fault Tolerant Power Compliance specification. (The first was the HP Server rp8400.)

## requirements for fault-tolerant power compliance

The rp7405 hardware is capable of receiving input from two different ac power sources. The objective is to maintain full equipment functionality, whether operating from both power source A and power source B or from either A or B alone. This capability is known as "fault-tolerant power compliance."

According to the Uptime Institute's Site Uptime Network specification, for equipment to qualify as being truly fault-tolerant-power-compliant, it must meet all of the following criteria as initially installed and as ultimately used in operation:

- If either one of the two ac power sources fails or is out of tolerance, the equipment is still able to start up or to continue uninterrupted operation with no loss of data or reduction in hardware functionality, performance, capacity, or cooling.
- After the return of either ac power source from a failed or out-of-tolerance condition during which acceptable power was continuously available from the other ac power source, the equipment does not require a power-down, IPL, or human intervention to restore data, hardware functionality, performance, or capacity.
- The first or second ac power source can fail one second after the return of the first or second ac power source from a lost or out-of-tolerance condition with no loss of data, hardware functionality, performance, capacity, or cooling.
- The two ac power sources can be out of synchronization with each other, with different voltages, frequencies, phase rotations, and phase angles, as long as the power characteristics for each separate ac source remain within the range of the manufacturer's published specifications and tolerances.

- Both ac power inputs terminate within the manufacturer’s equipment. Internal or external active input switching devices (such as static transfer switches) are not acceptable.
- A fault inside the manufacturer’s equipment that results in the failure of one ac power source will not be transferred to the second ac power source, causing it to also fail.
- An internal uninterruptible power supply (UPS), internal power batteries (batteries for cache memory are acceptable), or other type of energy storage equivalent is allowable only for the purpose of a prompt, orderly shutdown. The existence and volt-ampere capacity of an internal UPS or batteries, and the time required for a prompt orderly shutdown, must be identified.
- With both AC power inputs available, the power provided by each of the two internal power trains is 50% ± 10% of the power output for the supply.
- An external software alarm is provided via the equipment’s software or the host’s operating system when an ac power source is lost or is outside the manufacturer’s published tolerances, and the software also indicates when the abnormal condition is corrected.

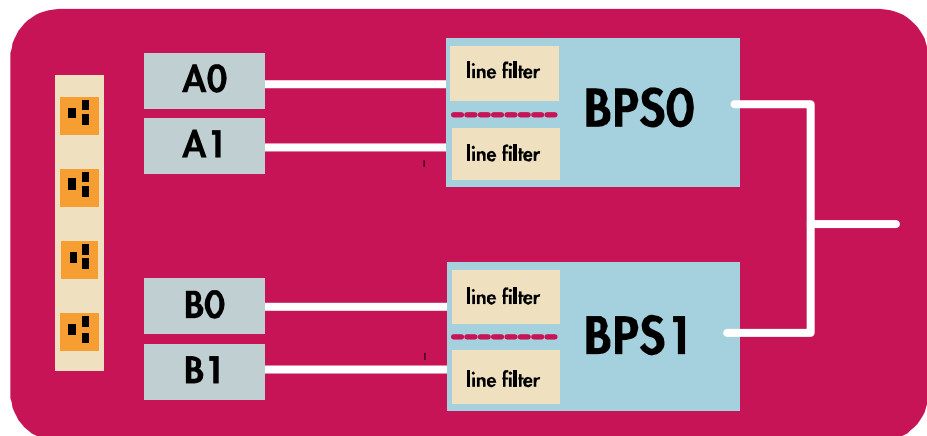
The HP Server rp7405 meets these criteria with no deviations.

**fault-tolerant power compliance in the rp7405**

The ac input to the rp7405 has been divided into four separate circuits. Each circuit can be fed by any 50 to 60 Hz source through four line cords. A minimum of two power cords must be used to maintain normal operation of the rp7405. A second set of two cords can be added to improve system availability by protecting, for example, against power grid failures, failed power supplies, or accidentally tripped circuit breakers. Four power cords must be used in order to enable redundancy and hot-swap functionality of the bulk power supplies. These four lines are labeled A0, B0, A1, and B1 at the line filter inputs on the back panel of the rp7405.

**Figure 12** shows the rp7405 power configuration. Note that each ac inlet feeds one rp7405 bulk power supply (BPS) and that two separate ac inlets feed each bulk supply. The design of the individual BPSs and the configuration of the interconnect meet all the conditions listed above for fault-tolerant power compliance. When using only two power cords, best availability is achieved by plugging one cord into BPS0 and the other into BPS1.

**figure 12. Power inputs and interconnects in the rp7405 are designed for fault-tolerant power compliance.**



## features of the ac power system

The rp7405 ac power configuration has several salient features designed for the convenience of the user. First, and most important, are the four 16-/20-ampere power cords used for 2N power. At first glance, it may appear that using two higher-amperage cords would have been better. HP considered this solution, but rejected it because this would have required either two 20-ampere 3-phase cords or two 30-ampere single-phase cords. Research has shown that many customers do not have three-phase power in their data centers. In short, using four lower-amperage cords is the only solution that will fit all data centers.

## 20-ampere circuits

In the rp7405, the 20-ampere circuits provide:

- Flexible, manageable, small-radius cables to and inside the cabinet containing the rp7405 system
- Economical cables, plugs, receptacles, and wiring that keep costs down
- Easy system installation because of the general availability of 20-amp circuits
- Manageable cabinet internal wiring, coupled with the ability to use just two single high-amperage cables for mains power to the cabinet, via the addition of large-amperage cabinet-housed power distribution units (PDUs)

## dual ac modular bulk power supplies (BPS)

Each BPS shown in figure 12 can be thought of as two sub-power supplies housed in a single BPS module. If only Grid 0 power cords are used, only one of the two sub-power supplies is in use. The other supply sits idle until a second set of Grid 1 cords are used. The benefit of this design is that when utilizing all four power cords, the rp7405 has 2N redundant power protection.

The dual ac modular bulk power supplies (BPSs) provide:

- redundancy for both hardware failures and power input failures
- hot-plug capability for any BPS in a redundant configuration
- better data security, maintenance scheduling, and maintenance operations without system interruption

## ac power consumption

The HP Server rp7405 power consumption varies greatly depending on the hardware configuration and the input line voltages supplied at the customer site. Because energy efficiency is of such high importance all over the world, HP represents power consumption in terms of voltamperes (VA). There are two ways to represent consumption:

- **theoretical maximum value**—represents the maximum wattage of a given configuration, assuming worst-case conditions (thermal tolerances, workloads, and so forth) on all system components. It is extremely unlikely that any customer will experience power consumption at the level specified as the theoretical maximum value.
- **typical value**—represents the expected power consumption of a given configuration. The typical value is the approximate power consumption that a customer will most likely experience and can use for power budgeting purposes.

## maximum and typical power consumption values

Following are the theoretical maximum and typical power consumption values for two different rp7405 configurations:

### fully loaded rp7405 configuration

The rp7405 fully loaded configuration consists of eight 650-MHz PA-8700 processors, 32 GB of memory, 14 PCI cards, two cell boards, four internal hard drives, one DVD drive, two core I/O card sets, and two bulk power supplies.

- Theoretical maximum power consumption: 3,000 VA (15A @ 200 Vac)
- Typical power consumption: 1,700 VA (8.5A @ 200 Vac)

## 60-ampere power distribution unit

### average rp7405 configuration

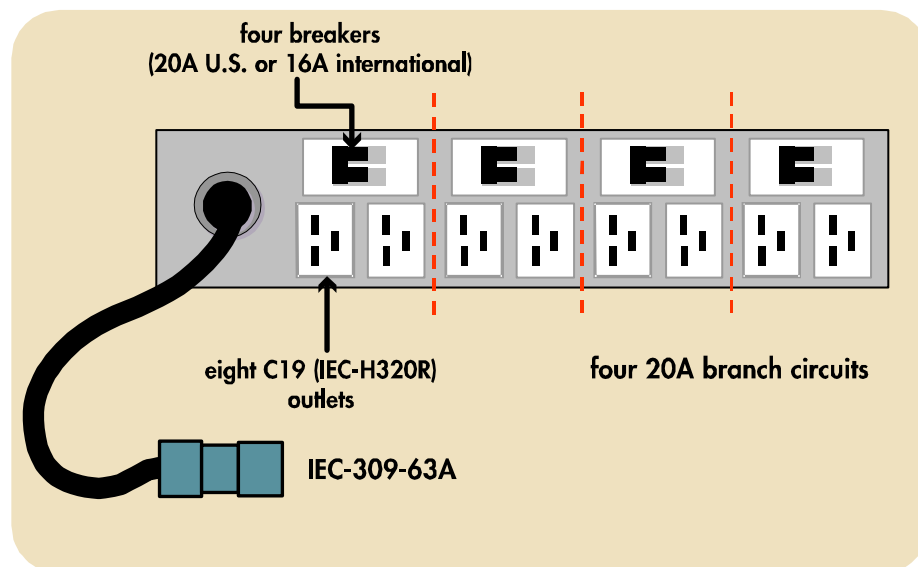
The rp7405 average configuration includes four 650-MHz PA-8700 processors, 10 GB of memory, five PCI cards, two cell boards, two internal hard drives, one DVD drive, one core I/O card set, and two bulk power supplies.

- Theoretical maximum power consumption: 2,150 VA (10.75A @ 200 Vac)
- Typical power consumption: 1,050 VA (5.25A @ 200 Vac)

Customers who prefer to use a power distribution unit (PDU) instead of directly plugging into wall outlets have three PDU solutions to choose from. Use of a PDU offers the benefit of adding an additional breaker into the power scheme and reduces the number of required wall outlets.

Customers who prefer the fewest higher-amperage connections from their ac line current source to the rp7405 can use the 60-ampere power distribution unit (PDU). This PDU is designed to fit horizontally in a standard 19-inch cabinet and occupies three EIA units of racking space. The 60-ampere PDU is sold separately and can be ordered with any HP server solution. Product numbers are E7683A (North America) and E7684A international). **Figure 13** shows the 60-ampere PDU.

**figure 13.** The 60-ampere PDU provides four breaker-protected outlets from a single ac line input.



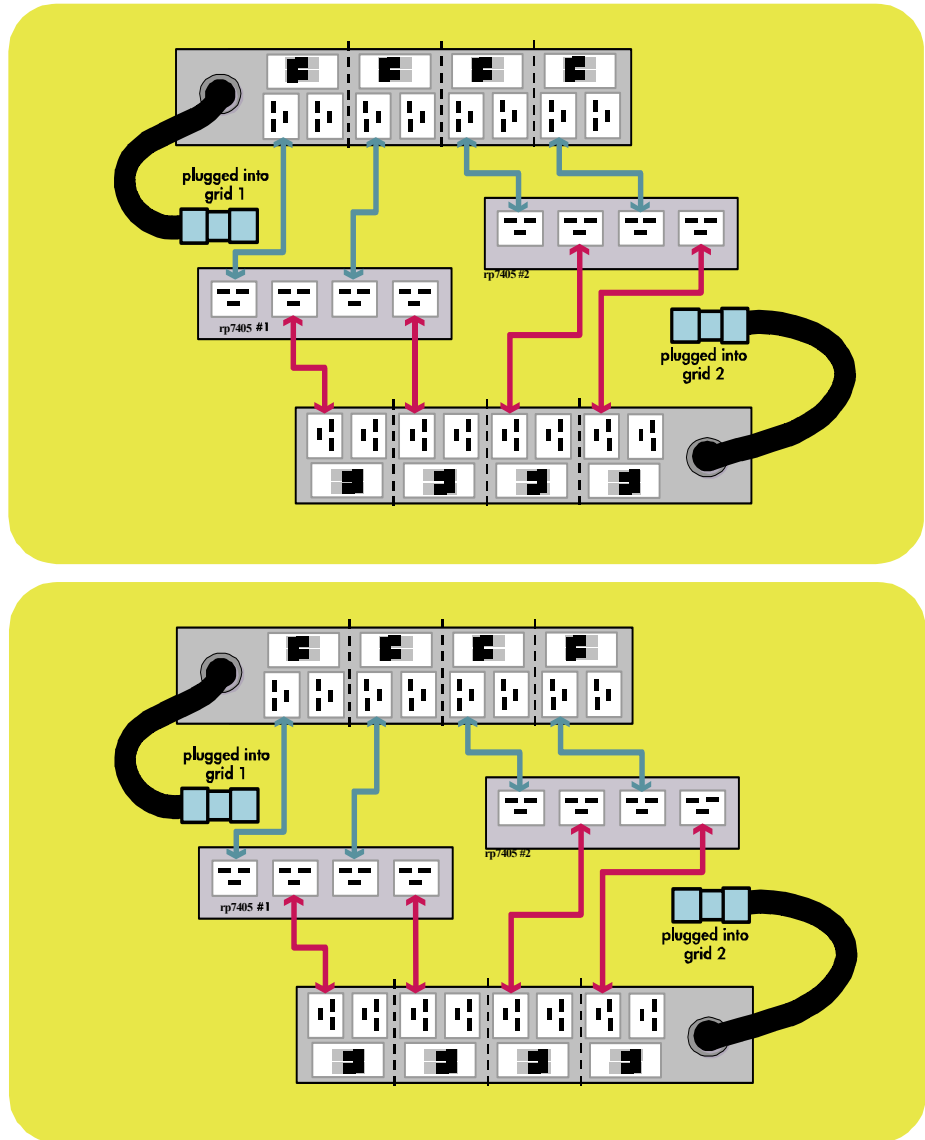
Each 60-ampere PDU contains eight C19 outlets spread evenly among four 20-ampere branch circuits (two C19s per branch). Each of the four branch circuits is protected by a circuit breaker that is either 20 amperes (United States) or 16 amperes (international). All 60-ampere PDUs are delivered with an IEC-309-63A plug.

The maximum amperage is 60 amperes through the entire PDU and 20 amperes per breaker. Both limits must be met. If 20 amperes are being drawn per breaker, only three sets may be used.

Each 60-ampere PDU can support up to four rp7405 servers if the PDU is not mounted in the rack. Up to three rp7405 servers can be supported if the PDU is mounted within the rack. (Each 60-ampere PDU consumes three units of rack space.)

**Figure 14** illustrates an example of using 60-ampere PDUs to configure redundant power for two rp7405 servers.

**figure 14. Configuring two rp7405 servers with two 60-amp PDUs for redundant power.**



The blue cords (that is, the top four double-headed lines) are the primary power connections needed for normal operation. In this example, both primary connections from each server are plugged into separate branch circuits. The remaining PDU outlets can be used to power other components, as long as the rating of 16/20 amperes per PDU branch circuit breaker and 60 amperes per PDU is not exceeded.


For redundant power inputs, the red cords (the bottom four double-headed lines) are added. These represent secondary power connections. If (as shown in the figure) the second PDU is plugged into a second grid, this configuration provides protection against:

- Losing power from a single power grid
- Accidental tripping of one or two circuit breakers
- Accidental disconnection of a single PDU power cord
- Accidental disconnection of up to two power cords per system
- Failure of a single bulk power supply (BPS)

## 30-ampere power distribution unit

Rack configurations consisting of peripherals and only one HP Server rp7405 will likely be best supported with the 30-ampere PDU. This PDU is also designed to fit horizontally in a standard 19-inch cabinet but occupies just 2 EIA units of racking space. The PDU is sold separately and can be ordered with any HP server solution. The 30-ampere PDU's product numbers are A5499AZ, option 001 (North America) and A5499AZ, option 002 (international).

**figure 15.** The 30-ampere PDU provides two breaker-protected branch circuits from a single ac line input.




Each 30-ampere PDU contains eight C13 and two C19 outlets spread evenly across two 20-ampere branch circuits (one C19 per branch). Each branch circuit is protected by a circuit breaker that is either 20 amperes (United States) or 16 amperes (international). All 30-ampere PDUs are delivered with an IEC-309-30A plug.

The maximum amperage is 30 amperes through the entire PDU and 20 amperes per breaker. Both limits must be met.

Each 30-ampere PDU will support two cords from a single HP Server rp7405, with additional C13 outlets available for peripherals. When all four of the rp7405 cords are used, two 30-ampere PDUs must be used. (There are two C19 outlets per PDU.)

**figure 16. Example of using 30-ampere PDUs to configure redundant power for one HP Server rp7405.**



The blue cords (that is, the top two double-headed lines) are the primary power connections needed for normal operation. In this example, both primary connections from each server are plugged into separate branch circuits. The remaining PDU outlets can be used to power other components, as long as the rating of 16/20 amperes per PDU branch circuit breaker and 30 amperes per PDU is not exceeded.


For redundant power inputs and to enable hot-swap BPS functionality, the red cords (the bottom two double-headed lines) are added. These represent secondary power connections. If (as shown in the figure) the second PDU is plugged into a second grid, this configuration provides protection against:

- Loss of power from a single power grid
- Accidental tripping of one or two circuit breakers
- Accidental disconnection of a single PDU power cord
- Accidental disconnection of up to two power cords per system
- Failure of a single bulk power supply (BPS)

#### 16-ampere power distribution unit

There are two power options available for rack configurations with no available EIA space for PDUs. The first option is to plug the rp7405 cords directly into the wall outlets. The second option is to use 16-ampere PDU strips. These 16-ampere PDUs can be mounted in either horizontal or vertical positions within an HP rack. In the vertical orientation these PDUs don't consume EIA space, and up to six can be mounted into a single 2-meter rack. This PDU is sold separately and can be ordered with any HP server solution. The 16-ampere PDU's product number is E7671AZ.

figure 17. The slim 16-ampere PDU won't consume valuable EIA rack space.



Each 16-ampere PDU contains six C13 and two C19 outlets. These outlets are protected by a single 16-ampere circuit breaker.

Each 16-ampere PDU will support two cords from a single HP Server rp7405, with additional C13 outlets available for peripherals. Extra care must be taken to not overload the PDU when adding peripherals. When all four of the rp7405 cords are used, two 16-ampere PDUs must be used. (There are only two C19 outlets per PDU.)

## performance and scalability

The HP Server rp7405 maintains leading-edge performance and scalability. The compact rp7405 offers performance density at least double that of the competition. For example, four rp7405 servers fit into a single HP rack. Competitive systems, on the other hand, require at least two cabinets, or even four, to equal the performance of this rp7405 configuration.

## I/O bandwidth

The HP Server rp7405 features excellent I/O and memory bus bandwidth:

Cell-controller-to-I/O-subsystem bandwidth (peak)	4 GB/s
I/O slot bandwidth (peak)	8.5 GB/s
Memory bus bandwidth (peak)	8 GB/s

## scalability

The HP Server rp7405 also has scalability that allows flexibility in configuring the system for present and future requirements:

feature	capacity
Cell boards	1–2
CPUs	2–8
Memory (with 512 MB DIMMs)	1–32 GB
Hot-plug PCI I/O slots	15 slots (64-bit x 66 MHz)
Partitions	1 or 2
Hot-plug internal disks	0–4 (18, 36, or 73 GB disks)
Hot-plug removable-media drives	0–1 (DVD included or DAT)

## performance considerations and configuration rules

Applying a couple of simple configuration rules will ensure top performance with the rp7405:

configuration rule	why it's applied	details
<b>A cell should be configured with a multiple of 8 DIMMs.</b>	Maximizes cell memory bandwidth	A minimum of 8 DIMMs are recommended in order to populate both memory buses on a cell board. There may be additional interleaving and/or bandwidth benefits to populating more rows (16 DIMMs). With only 4 DIMMs, only ½ of the memory bandwidth will be utilized.
<b>All cells in a partition should have the same amount of memory (symmetric memory loading).</b>	Best memory interleaving	Asymmetrically distributed memory affects the interleaving of cache lines across the cells. Asymmetrically distributed memory can create memory regions that are non-optimally interleaved. Applications whose memory pages land in memory interleaved across just one cell can see up to four times less bandwidth than applications whose pages are interleaved across all cells.

## **racking**

The HP Server rp7405 provides industry-leading performance density and availability in a racked configuration. At 10 EIA units (17.5 inches) each, four rp7405 servers can be mounted into a single HP two-meter cabinet.

The rp7405's industrial design and packaging allow easy and quick access to all of the system's components. The most frequently handled removable-media devices and disks are directly accessible at the system's front. By removing the front bezel, hot-swap fans, hot-swap power supplies, and PCI power supplies can be completely serviced. At the rear, core I/O components and additional hot-swap fans are directly accessible.

For access to all other components, the HP Server rp7405 comes with rack slides. These enable the rp7405 to slide forward out of the cabinet for servicing of internal components such as fans, cell boards, and I/O cards—even while the system is still operating.

The slides also allow for servicing or replacement of any FRU (field-replaceable unit) without removing the chassis from the cabinet. In fact, any FRU can be accessed and removed within 15 minutes or less. This design minimizes the downtime associated with system upgrades in the rare event of a component failure.

Also included with every rp7405 is a cable management arm (CMA). The CMA neatly secures data cables and prevents them from becoming entangled while the system is being serviced.

## **ballast kits**

For stability during servicing, ballast kits are used on HP cabinets in which the rp7405 is installed. Every rp7405 shipped to customers includes a ballast kit (the same ballast kit used for the rp7410 and rp8400). The ballast easily attaches to the rear anti-tip foot that comes standard with every HP Rack System E cabinet. Use of the rp7405 ballast kit is mandatory, and the kit must be installed immediately.

## **third-party racks**

HP servers are designed to maximize performance density when installed into HP system cabinets. Moreover, HP system cabinets maintain the high level of safety and reliability that customers have come to expect. Although HP strongly recommends racking the rp7405 in HP cabinets, we recognize that some customer circumstances may prohibit this. So HP has developed guidelines that enable safe, reliable HP server installations in third-party cabinets. Because of the wide variety of cabinets in the marketplace, it is extremely important that the guidelines be followed explicitly. (See chapter 4 in the *HP Server Configuration Guide* for the racking guidelines.)

## nPartitions

The HP Server rp7405 offers the same mainframe-style partitioning capability introduced in the HP Superdome, rp8400, and rp7410 servers. When coupled with the HP-UX operating environment, partitions in the rp7405 allow a single system to be logically divided into two independent systems within the same server.

Partitions in the rp7405 are referred to as nPartitions; the “n” indicates that a partition actually has a wide number of functions that it can perform. These nPartitions provide a high degree of isolation that can be utilized to consolidate several applications on a single server, give flexibility in resource management, and deliver improved availability and scalability.

## how partitions work

A hardware partition corresponds roughly to a single, standalone system. The rp7405 can be subdivided into two partitions.

Special programmable hardware in the cells defines the boundaries of a partition in such a way that isolation from the actions of other partitions is enforced. Each partition runs its own independent instance of the operating system (at first release, HP-UX 11i only). Non-clustered applications cannot span partitions, since each partition runs its own instance of the OS—essentially functioning as a standalone server. However, different partitions may be executing the same or different revisions of an operating system. And at a future date, when Intel® Itanium™ Processor Family processors are available on the rp7405, different partitions will be able to execute different operating systems—such as HP-UX, 64-bit Windows®, and Linux—any two at one time.

Each partition has its own independent CPUs, memory, and I/O resources. With a future release of HP-UX, using the related capabilities of dynamic reconfiguration (for example, online addition and online removal), administrators will be able to add new resources to a partition or remove and replace failed modules while the partition continues in operation.

## nPartition configurations

As previously mentioned, the rp7405 can be divided into two independent hardware partitions. In a partitioned configuration, system resources such as I/O slots, core I/O, and disks are evenly split between the two partitions. There is no flexibility to otherwise divide these components. For example, it is not possible to include 12 I/O slots in partition 0 and four I/O slots in partition 1; the split must be even. The only exception is the removable-media bay, which is only attached to partition 1.

## reconfiguring partitions

Increasing or reducing the processing power of a partition is accomplished, respectively, by adding a cell to a one-cell system or deleting a cell from a two-cell system. (For the first release of the rp7405 with HP-UX 11i, the operating system running on the affected partitions must be shut down before cells are moved and before configuration changes will take effect.) Although HP-UX 11i does include commands for some configuration tasks, HP recommends using the Partition Manager (parmgr) to configure partitions.

At first release, the combination of the rp7405 and HP-UX 11i supports hardware partitioning. Hardware-based partition configuration changes may require a reboot of the partition, depending upon the configuration change. The reboot of the partition only affects the partition that is being reconfigured; the other partition is not affected and will continue to execute without interruption.

## dynamic reconfiguration

In a future HP-UX release, dynamic hard partitions will be supported. Dynamic partitions mean that partition configuration changes do not require a reboot of the partition. Virtual partitioning within HP-UX 11i also provides additional options for dynamic reconfiguration without having to reboot the changed partition.

## uses for nPartitions

One of the primary advantages of nPartitions is their support for flexible, user-defined partitioning of resources. They also isolate hardware and software errors. Here are some of the typical types of nPartitions:

- **OS partitions**—A partition can run various operating environments such as HP-UX, 64-bit Windows, or Linux. (Only HP-UX is available at first release.)
- **OS version and functional partitions**—Partitions on a system can run different versions of the same operating system, and each can have a specific function (for example, a database partition, an application partition, a development partition, a production partition, and so on).
- **high-availability partitions**—Partitions can participate in an MC/Serviceguard cluster, with applications failing-over from one partition to another.
- **application partitions**—A partition hosting a specific application can provide dedicated resources, such as memory space.
- **floating partitions**—Applications will be able to take advantage of policy-based workload balancing across partitions (not supported at first release).
- **virtual partitions**—Partitions can support workload consolidation and software isolation with better granularity (for example, to the individual CPU level).

## more details about some key uses for nPartitions

Partitioning resources provides extraordinary flexibility in configuration and operation. Partitions allow allocation of physical resources and tuning of the operating system running on each partition, depending on the needs of the application (or the most important application) being run. Alternatively, the rp7405 can be configured as a single partition, allowing all the resources to be focused on a single set of tasks—for example, a large online transaction-processing application.

This section gives more details about some of the most popular and effective uses for nPartitions.

## resource partitioning

Today, businesses require that multiple applications running on the same server deliver full performance and do not conflict or interfere with one another. For top performance, it is critical that the actions of one application not adversely impact or result in the denial of service to other applications. Furthermore, a failure in the operating system or hardware components on the server supporting a particular set of applications must not result in the loss of multiple applications running on other partitions on the same server.

In order to have effective consolidation of applications on a single server, an application must not be allowed to use a disproportionate share of system resources needed by another application, and the failure of a hardware component must not result in the loss of service for multiple applications. While a solution to these problems could be accomplished by providing multiple servers at a site, such a method is inflexible in the face of changing requirements, is more expensive, and requires more space in the data center. Instead, the IT department needs a flexible configuration that offers the capability to adjust compute resources in a dynamic fashion while still preserving the SMP programming model.

That's why nPartitions are so effective at resource partitioning. With rp7405 nPartitions, each partition is isolated from the other with special, dynamically configurable hardware that provides an effective firewall between applications. There is no possibility of interference between nPartitions.

## multiple environments

Business requirements and competitive pressures can dictate the need for a large set of applications—both old and new. Often applications require different versions or revisions of operating systems, different system parameter settings, and even different patch levels. And performance can suffer from the interactions of applications that have incompatible behaviors.

But nPartitions can solve the problem of different versions or revisions by running multiple environments. With these partitions, a portion of the compute and I/O resources can be configured to run an old copy of the operating system while the rest of the machine is updated to run the newer software. Such a scheme can also be used to set up a test environment where new software and operating systems can be installed, leaving production applications to continue operations without being affected.

Similarly, two rp7405 servers can be set up such that each one runs the prime partition of an MC/Serviceguard cluster as well as the backup partition of the other's prime partition. This eliminates the few single points of failure between the prime and backup partitions.

## improved availability

Another use for nPartitions is to provide independence of failure or operation. If one partition must be rebooted or taken down to perform an operating system upgrade, the other partition in the system is not affected.

Unlike a traditional large SMP server, nPartitions allow a server to be configured into a cluster of independent systems. Since failure independence is provided by nPartitions, it is possible to configure high availability clusters within the same complex, with each cluster consisting of two or more partitions. In order to achieve maximum availability, the partitions making up the cluster can be configured as high-availability partitions on different nodes. This ensures that no single hardware or software failure can affect another cluster member, yet it retains the configuration flexibility of a single pool of system resources.

## more flexible scaling

Some applications may not have the ability to scale linearly with large numbers of CPUs. However, nPartitions allow the amount of processor, memory, and I/O resources to be tailored to match application-scaling characteristics. Moreover, as new application or operating system releases provide improved scaling characteristics, nPartitions allow resources to be readjusted to optimize application performance.

## what you gain with nPartitions

To sum up, the major benefits of nPartitions are:

- Increased server manageability and flexibility
- Reduced total cost of ownership
- Increased availability for mission-critical applications

Other benefits include:

- Independent resource controls, I/O, and security in each partition
- Ability to run multiple applications and/or different operating systems on a single system, enabling the best use of computing resources while providing hardware and software fault and security isolation
- Protection against software errors—a software error in one partition (including an OS failure) will not bring down other partitions within the server
- Protection against hardware errors in other nodes—a hardware error in a separate coherent node will not bring down a partition in any other node
- Protection against hardware errors in the same node—most hardware errors in the same node will not bring down a peer partition in the same node
- Protection against downtime due to configuration changes in another partition—configuration changes in a different partition (such as online replacement or addition of a component, reboot, shutdown, and so on) will not affect other active partitions on the same or a different node

## high availability

High availability (HA) is the hallmark of HP computer systems. But HP knows that delivering solutions that fully enable the “always-on,” 24 x 7 operations demanded of today’s businesses requires more than just delivering laundry lists of unusable HA features—or HA features with limited utility. The high-availability features of the HP Server rp7405 actually address the *real* causes of customer downtime, as determined by actual field data from midrange computer users.

## partition reliability

The rp7405 has a design that is significantly “hardened” over other systems in its class. In fact, many of the features in this system can only be found in mainframes (or HP rp8400, rp7410, and Superdome systems). The availability features within each rp7405 partition have been field-proven to provide high system availability. And many customers who have taken advantage of these features report significantly lower hardware failure rates than with competitive systems.

## CPU protection

The central processing unit is often a major cause of system downtime. For instance, CPU cache errors are demonstrated to be a large contributor (in many cases, the greatest contributor) to unplanned system downtime. Furthermore, addition or modification of CPU resources is among the highest-ranking causes of planned hardware downtime. But in the rp7405, HP has designed specific features to combat CPU-caused downtime, including:

- Full error checking and correction (ECC) on all caches
- Automatic deconfiguration of “faulty” CPUs—known as dynamic processor resilience (DPR)
- A highly effective and reliable CPU cooling scheme
- Redundant CPU power converters

### **ECC on caches**

The CPU caches in the rp7405 are fully protected from single-bit hard errors and random soft errors generated from cosmic rays or other intermittent error generation sources. Some competitive systems in the same class are not similarly protected, resulting in errors that are hard to debug and that in many cases are blamed on the customer environment. Such cache errors in these unprotected systems can result in failures that bring down multiple partitions.

Another advantage of the rp7405’s CPU cache is its layout, which significantly reduces the chance of a multi-bit error due to a random cosmic ray strike. Such attention to detail is not found in many designs available from other vendors.

### **automatic CPU deconfiguration**

Dynamic processor resilience (DPR) refers to the ability of the system to detect and de-allocate CPUs that are generating excessive (recoverable) single-bit cache errors. This prevents data-corrupting double-bit cache errors that can result when single-bit errors become abundant. By de-allocating offending CPUs before extreme levels of single-bit cache errors are generated, double-bit cache errors are all but eliminated. DPR provides the customer with assured data cache integrity, preventing hard errors from occurring and causing downtime.

Here's how DPR works:

1. Processor detects single-bit error in data cache and vectors to processor-dependent code (PDC).
2. PDC generates a low-priority machine check (LPMC).
3. LPMC handler logs information to diag2 driver.
4. Diaglogd daemon pulls LPMC log information from diag2 and passes it to the HP Event Monitoring Service (EMS) LPMC monitor.
5. If there have been too many LPMCs within the preceding 24 hours, CPU is de-allocated (online).
6. System firmware is called to have PDC disable the processor the next time the system boots.
7. Event is generated to notify customer and HP.

This functionality is currently available for all CPUs in a partition except for the Monarch CPU. The Monarch processor refers to one processor that is selected during system boot and given special boot and interrupt responsibilities. Although the Monarch CPU will continue to correct cache errors "on the fly," it is not de-allocated until the next reboot. This limitation will be eliminated in a future operating system release.

### **system cooling**

Heat is the big enemy of electronic components. But the rp7405's two-level cooling scheme offers outstanding cooling capacity at a nominal cost. High-efficiency system fans provide redundant enclosure cooling for all operating environments and system loads. This is the only level of cooling provided on most competitive systems. HP goes on to provide additional local cooling directly at the hottest system components. Turbo-cooler fans draw air directly into the heat sinks of the CPU and cell VLSI. At the extremes of the rp7405's operating range, the turbo-cooler fans keep temperatures well below the maximum values allowed. Even though the turbo-coolers may not be required under normal operating conditions, running them ensures the silicon chips operate at the lowest possible temperature, ensuring maximum component life and reliability.

To further improve the availability of the rp7405, manageability software monitors the speeds of all fans, including turbo-cooler fans. The rp7405 Smartfan controller can detect the first hint of slowdown associated with bearing wear, enabling the customer to replace marginal cooling components at the next scheduled maintenance interval.

### **redundant CPU power**

In the rp7405, CPU power is protected by means of component redundancy. Power to each of the CPUs is provided from a pool of N+1 dc-dc power converters. No single converter in this pool can cause system downtime.

## memory protection

Main memory failures are the single largest cause of customer downtime. The rp7405 has several features designed to reduce or eliminate failures of memory:

- “chip kill” tolerance
- Dynamic memory resiliency (DMR)
- Automatic deconfigure on reboot
- Hardware memory scrubbing
- ECC protection

### **chip kill tolerance**

“Chip kill” tolerance is the ability of the system to continue to run in the face of any single- or multi-bit chip error on a DRAM chip. The DRAM chips in the rp7405 are N+1 redundant per set of 128 DRAM chips per memory word. This functionality is essential in the design of reliable memory systems, and systems without this feature are significantly more susceptible to failure compared to the rp7405. (This has been demonstrated at customer sites that use both chip kill tolerance and less reliable architectures.)

In systems where hundreds of high-speed DRAM chips are employed in a memory structure, the likelihood of isolated errors becomes significant. Customers depend on the performance and integrity of every active memory segment in order to maintain system coherency. Only through HP’s “chip kill” memory redundancy are these levels of availability, scope, and performance assured.

### **dynamic memory resiliency**

Dynamic memory resiliency is the system’s ability to de-allocate failed memory pages online. This feature is similar to dynamic processor resiliency; if a location in memory proves to be questionable (that is, exhibits persistent single-bit errors), the memory is de-allocated online with no customer-visible impact. Assuming the rp7405 is equipped with adequate memory to begin with, it is likely that the failed memory will never have to be replaced over the life of the product, resulting in a significant reduction in both planned and unplanned downtime.

## protection for I/O and disks

I/O errors are another significant cause of hardware errors and downtime because:

- The number of I/O cards in a typical system is significant
- The I/O cards themselves are a part of the system most exposed to frequent human interaction in the data center

In order to prevent downtime due to I/O errors, HP has designed the following features into the rp7405:

- Online replacement of PCI cards
- Hardware “firewall” of I/O errors to the cell level
- High-MTBF (mean time between failures) I/O cards
- Separate PCI buses for each I/O card

Taken together, these features will reduce hardware downtime by at least 20 percent over similar servers.

### **online replacement of PCI cards**

PCI cards in the rp7405 can be added and replaced online at product introduction, with delete capability planned for a future HP-UX release. Online addition/replacement reduces planned downtime by allowing the online addition of resources to a running system and the replacement of a failed card without bringing down the system.

### **hardware “firewall” to limit I/O errors**

No I/O card in the rp7405 shares a PCI bus with any other card. In other words, all I/O cards are logically and electrically isolated from all other PCI cards. This significantly improves performance, reduces the likelihood of I/O errors, and confines errors to a single PCI bus and card.

### **high-MTBF I/O cards**

HP’s I/O cards are extensively qualified to work in our servers. This means stringent MTBF goals and full interoperability testing of all I/O cards: no incompatible drivers, no bad pointers, and no marginal signal quality.

### **disk drive redundancy**

For critical applications the boot drives in the rp7405 can be mirrored. In single-partition systems, this enables failure resiliency of the disk mechanism, SCSI controller chips, and bus bridges. HP offers deeper hardware redundancy in the SCSI strings than most competing brands. In dual partitions, both partitions can provide boot mirroring, but only one partition will have deep controller and bus redundancy. This level of hardware enablement, along with efficient software support, makes online disk deletion and replacement transparent to system operation. In this way system availability is assured.

## **reliability in the cabinet infrastructure**

In keeping with its focus on maintaining high availability (HA), the rp7405 includes protection against failure within the cabinet infrastructure. The HA features in this area include:

- True dual ac line cord support (phase independent)
- 2N redundant system power supplies
- Complete resilience to service processor failures

### **dual ac line cord support**

As described earlier in this paper, the rp7405 server can run on one or two totally independent power sources. Moreover, these two power sources do not need to be in phase—or even at the same frequency.

### **redundant power supplies**

Two identical system power supplies are included in the rp7405 base system. Either supply is fully capable of powering a maximally configured system. Normally the two supplies equally share system energy needs to minimize load stress on both supplies and enhance reliability. In the event of a supply failure, the remaining supply enables full system capability, eliminating downtime due to supply failure.

### **resilience to service processor failures**

The rp7405 hardware has been designed to enable service processor failover when redundant core I/O cards are in place. Future firmware and manageability code releases will allow a slave service processor to take over for a failed master, and will also enable a resilient console (on reboot). Future OS releases may allow the console to fail-over as well.

## serviceability

HP continues to make great strides in implementing features that reduce the time to upgrade components or diagnose and repair component failures. The HP Server rp7405 was designed with the objective of coupling state-of-the-art diagnostic tools with hardware features to virtually eliminate unplanned downtime. Hot-swap/-plug technology is implemented throughout the server, allowing addition or replacement of components while the system continues to run.

For components that cannot be serviced while the server is running, the rp7405 was designed to provide access and removal of any field-replaceable unit (FRU) within 15 minutes or less. In addition, the HP Server rp7405 is loaded with design innovations that greatly simplify servicing.

## enhanced serviceability

The following is a list of rp7405 for enhanced serviceability features:

- Hot-plug functionality for internal disks, PCI cards, and removable-media drives (also, dynamic hot-plug of cell boards will be enabled with a future release of HP-UX)
- Hot-swap functionality for bulk power supplies and cooling fans
- Product depth from mounting column to the connector-mating surface at rear of 26.75 inches (~30 inches overall with the front bezel sitting forward of the rack column), leaving 4 inches for cable bending in third-party racks
- PCI card access from the top of the chassis so that ceiling light enhances visibility
- Access panels as large as possible to enhance the service access area
- Five-sided access to maximize accessibility while the product remains racked
- Independent access to major FRUs; any FRU can be accessed and removed within 15 minutes or less
- Access to most commonly serviced components from the front or rear without moving the package in the rack

## easy access to all components

Access on the top and sides requires the unit to slide out on slides. A lock in the slide latch prevents inadvertent sliding when the rp7405 is being moved.

The HP Server rp7405 extends 31 inches, allowing 3.5 inches of space for access to the external rear PCI cables, core I/O cables, and ac power cables. Once the chassis is extended on slides, access is achieved by removing the top and side covers held on with captive hardware.

For safety, an interlock device allows only one device to be extended at a time. A front foot and ballast system provide rack stability. (The rack ballast and front foot are required for safety.)

A large cable management arm, included with each server, facilitates cable dressing when the rp7405 is extended from the rack.

### front access

- Bulk power supplies
- Front system fans
- PCI dc-to-dc converters
- Hard drives
- DVD/tape drives (these devices are connected with a service loop located inside the media bay, allowing the drives to be serviced without the need for rear or side access)

#### rear access

- Rear system fans
- PCI card cables
- Power cables (ac)
- Core I/O modules and cables

#### right side access

- Cell boards (These are located on the right side because the front door of the HP RB11 rack hinges on the left. This allows the cell boards to slide out some 16 inches for removal. To shorten critical bus lengths, the cell boards are mounted upside down. To access RAM and CPUs, lay the cell board on a work surface.)
- PCI backplane (also requires top cover to be removed)
- PCI board-mounted voltage regulator modules (VRMs) (also requires removal of the cover and structural access panel)

#### left side access

- System backplane (a hinged backplane secured by jack screws makes backplane removal independent of the PCI FRU)

#### top access

- PCI cards
- PCI fan modules
- Mass storage backplane (held in place by two captive screws)

### attractive, functional design

HP's engineers recognize that industrial design and ergonomics play an important role in serviceability, and they have imbued the HP Server rp7405 with a high degree of design and refinement. These enhancements make the rp7405 both attractive and serviceable:

- Industrial design refinements on molded plastic rear fan modules
- Attention to industrial design on internal components such as the PCI OLR switches and paddles
- PCI fan modules molded in plastic
- Positive, audible snap sound for PCI fan module installation
- Internal plastic parts textured to enhance their cosmetic appearance
- Barcelona blue accent color to identify user access and service areas
- Colorful dividers for PCI cards that also provide a handle shape for improved grip
- Quick-release latching extractors on cell boards
- Smooth, snag-free ejector surfaces integrated into card guides
- Smooth, rolled sheet metal edges that eliminate possible skin injuries
- Captive hardware on all major FRUs
- Oversized slides that enhance stability and stoutness
- Graphics-rich labels and lots of label space (731 square inches) to provide information where it's needed
- Extensive use of LEDs for quick visual indication of system status and easy identification of failed components
- Serpentine ac cable retention bracket designed to accommodate a large range of cable diameters (anchors the ac power cable, occupies little space, and eliminates the need for cable ties)

## **ease of management**

As the number of servers grows and “server farms” proliferate, IT professionals have come to realize that the cost of managing these servers can add up to many times the actual cost of hardware. In the HP Server rp7405, HP has provided features to make management easier and less taxing. And with the addition of HP management software that interfaces seamlessly with the rp7405, system administrators are assured of a lower total cost of ownership (TCO) and higher overall efficiency.

## **rp7405 management processor**

The rp7405 has a dedicated management processor (MP) that simplifies and extends system management and enhances serviceability. The MP feature set was designed to minimize or eliminate the need for the system administrator to be physically at the system to perform tasks such as diagnostics, system management, or even hard resets.

See page 17 of this paper for a long list of management and control functions enabled by the rp7405’s management processor.

## **partition manager software**

The rp7405 comes with HP’s Partition Manager (parmgr) software, which makes specifying, setting up, and reconfiguring partitions easy and straightforward. Partition Manager can be launched as a GUI from HP’s System Administration Manager (SAM)—or directly from the command line. Partition Manager adds a number of valuable management capabilities, including:

- Display complex status
- Create and modify partitions
- Display a complete hardware inventory
- Display status of key complex components
- Check for problems or unusual complex conditions
- Manage power to cells and I/O chassis
- Turn on and off attention indicators for cells, I/O chassis, I/O cards, and cabinets

Partition Manager is another example of how HP makes server management easier than ever.

## **servicecontrol suite**

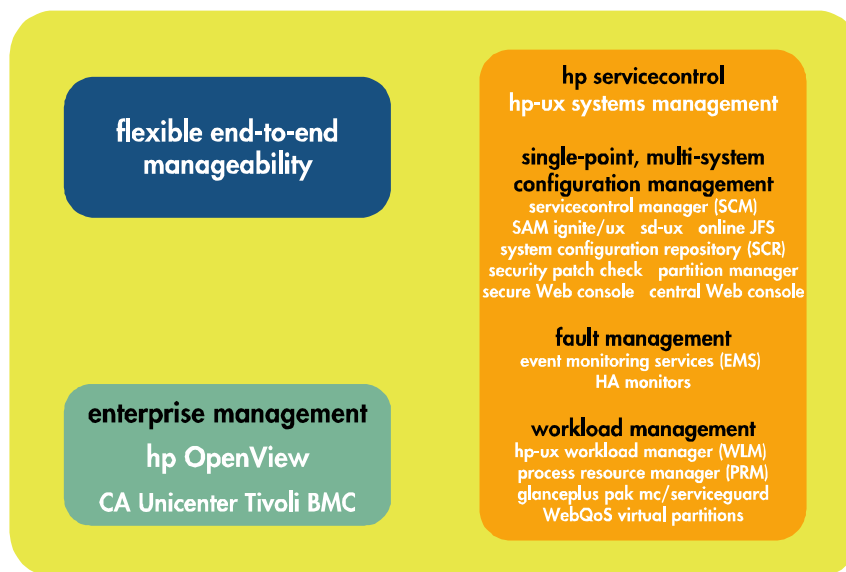
HP’s Servicecontrol suite provides increased efficiency for systems administrators through truly multi-system management tools. It provides flexible, end-to-end manageability for the entire server environment.

One important area addressed by the Servicecontrol suite is security. A concern in any IT environment—and especially important given the global reach of the Internet—security is addressed via capabilities such as role-based management and the highest degree of encryption.

Another valuable capability of the Servicecontrol suite is rolling out patches, updates, and new versions as well as keeping track of system usage, licenses, and assets. The Servicecontrol suite has provisions for rapid deployment as well as tools to enhance consistency and asset management. It also includes monitoring tools for keeping the administrator up-to-date on vital operating parameters such as performance, response time, and availability.

The Servicecontrol suite components are shown in **figure 18**.

figure 18. HP server Servicecontrol suite.



Servicecontrol suite ensures better control for the HP Server rp7405—and helps provide better and more profitable e-services.

## operating environment

As an operating environment, HP-UX 11i is a perfect complement to the HP Server rp7405. HP-UX 11i is HP's complete 64-bit UNIX® operating environment that delivers unparalleled scalability and performance for demanding applications. When teamed with HP's leading server systems, HP-UX 11i provides the power of supercomputing at a fraction of the cost.

## hp-ux 11i support

The rp7405 is supported by HP-UX 11i, March 2002 release or later. HP-UX 11i blends an extraordinarily robust kernel, blazing performance, and excellent scalability with a wealth of industry-leading add-ons from HP and third-party vendors. The result is an operating system that is simply without peer for the demands of end-to-end Internet-critical computing. HP-UX 11i is binary compatible with HP-UX 11.0 so there is no need to recompile. In fact, many 11.0 binaries run up to 20% faster on 11i.

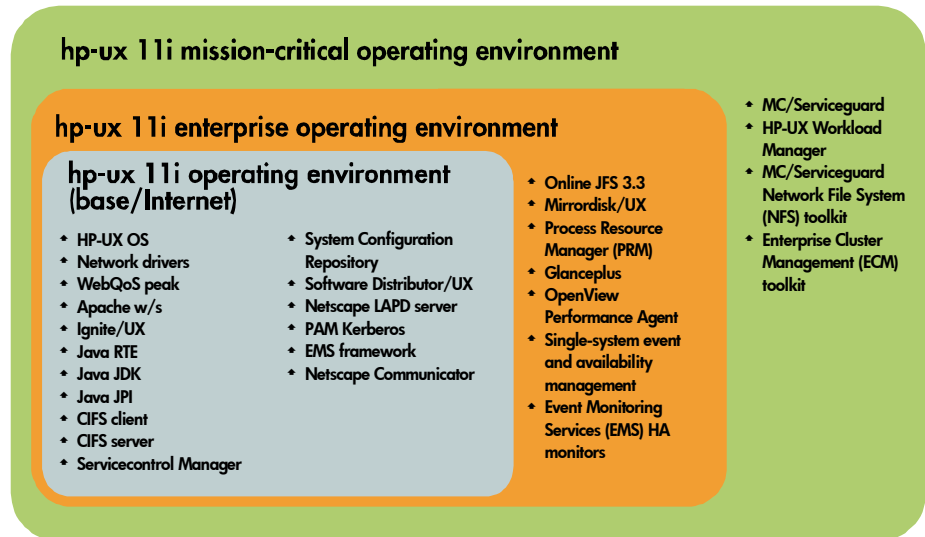
## key hp-ux 11i features

- **performance**—scalability up to 64-way multiprocessing, TPC benchmark disclosures, 4 GB large-process data space, threads support in FSS and 32-/64-bit DCE library
- **security**—host intrusion detection (IDS 9000), execute-protected stack to detect buffer overflow attacks, IPFilter host firewall. LDAP directory-enabled computing, Kerberos Server and PAM Kerberos, IPSec9000 VPN, HP-UX AAA Server and HP-UX Secure Shell, BIND 9.2.0
- **availability**—online addition and replacement of I/O and networking cards, dynamically tunable kernel parameters that take effect without reboot, up to 70% shorter system boot time
- **manageability**—partitions on rp7405 servers, desktop management interface (DTMI), Process Resource Manager physical memory control, Process Resource Manager disk bandwidth support
- **file systems**—VERITAS (JFS) 3.3, JFS ACL support
- **connectivity**—NFS over TCP/IP, Ethernet, Token-Ring, SNAplus, Frame relay and X.25, Cache FS, libc pre-enablement of DNS 8.X, libc IPv6
- **mass storage interface adapters**—includes adapters for SCSI, Fibre Channel, and RAID
- **internationalization**—EURO: ISO 8859-15 and 10646
- **graphics**—X11R6.3 server and runtime libraries
- ISU/ISV forward application compatibility from HP-UX 11.0

**choice of hp-ux operating environments**

HP’s engineers have packaged HP-UX 11i in a variety of *operating environments*, each aimed at a specific type of computing. These operating environments offer more focused functionality, easier ordering, and easier maintenance. HP-UX 11i operating environments install in half the time of Solaris 8 and a third of the time of AIX v5.1—with no codewords, so customers can be up and running in no time. **Figure 19** shows the contents of the various HP-UX 11i operating environments.

**figure 19. HP-UX 11i operating environments.**



**hp services for the hp server rp7405**

The success of an always-on computing environment depends on the right combination of technology products and services to ensure the infrastructure and processes meet the customer’s needs. HP services and support help customers derive real value from their IT investments through a broad services portfolio that includes consulting, outsourcing, support, financing, and education. HP’s industry-leading services portfolio offers a comprehensive set of services for servers that address all phases of a customer’s IT infrastructure lifecycle.

**design, migrate, and integrate**

HP offers a wide range of industry-leading services to help plan, design, integrate, and install a solution tailored to a customer’s business.

**consulting services**

HP consulting includes focused business integration, IT infrastructure, and specialized consulting services for global customers who need to design, integrate, and evolve a scalable environment. Utilizing best-in-class partnerships, we use proven methodologies, industry expertise, and IT experience to transform customer businesses.

**deployment services**

HP offers a full set of global implementation services—for HP and multivendor equipment—that allows customers’ IT solutions to be up and running, where, when, and how they want—anywhere in the world. Hardware and software integration can either be factory preconfigured or completed onsite, providing a custom solution that ensures a timely and successful installation. HP Project Management Services are also available to help manage the solution implementation.

hp startup integration express service	Startup integration is the highest level of integration service available from HP. This consulting service provides detailed design and integration, partitioning, and configuration of components to optimize performance and high availability. A certified HP consultant works with the customer to understand requirements and make technical specification recommendations for a custom-designed solution. This solution is integrated and ready "out-of-the-box" when it arrives at the customer site. "in"
hp smartest integration service	This service provides a preconfigured server that meets the customer's specific requirements. HP Smartset Integration is for customers who have the expertise to articulate partitioning requirements but would like to have the system delivered to them with the desired partitioning and hardware/software configuration already implemented.
hp onsite partitioning service	Since many customers don't want or need a partitioned server at initial shipment, this service allows the customer to conveniently implement a partitioned server environment in the field at some point after delivery and setup.
hp educational services	HP educational services can improve speed and productivity and help to keep IT staff up-to-date. HP offers comprehensive education solutions and content. We bring together the best strengths and expertise of HP and our partners and combine them with HP's learning approach to provide complete solutions that deliver focused learning on a one-time or ongoing basis.
<b>run and evolve</b>	HP offers the industry's broadest range of end-to-end global IT support services, providing proactive and reactive capabilities that flexibly scale to the demands of business-critical computing. HP support covers all assets and technologies, including hardware, software, networks, and even business processes.
standard support and industry-leading mission critical services	High-quality HP hardware and software support services, ranging from standard support to mission critical services, enable increased productivity, better availability of the computing environment, and improved return on the system investment. With standard support, the customer receives complete software maintenance and the flexibility to choose the hardware response time and coverage periods. Industry-leading mission critical services from HP are focused on proactively anticipating potential problems before they occur—and on providing fast resolution when problems do arise.
critical systems support	<p>To support the flexibility and scalability of the always-on Internet infrastructure, HP Critical Systems Support (CSS) is the optimal level of support. Designed for businesses that run critical applications and are severely affected by unplanned downtime, CSS provides technical expertise through an integrated combination of proactive services and fast problem resolution. HP helps customers meet availability and performance requirements through features that include:</p> <ul style="list-style-type: none"> <li>• An assigned account team</li> <li>• A detailed account-specific support plan</li> <li>• Change management assistance</li> <li>• Readiness assessments</li> <li>• A six-hour call-to-repair hardware commitment</li> </ul>

business continuity support	<p>HP Business Continuity Support (BCS) is the most comprehensive high availability support program HP offers. This service includes:</p> <ul style="list-style-type: none"> <li>• An assigned team of engineers</li> <li>• A customized plan of services specifically designed to eliminate downtime</li> <li>• Industry-leading tools and technologies geared to raise the productivity of the computing environment</li> </ul> <p>BCS customers receive top priority from HP laboratories, HP management, and HP response centers to ensure systems are performing as designed. BCS customers also receive preventive and proactive support, coupled with the fastest, most customized restoration commitment available today.</p>
support for data storage	<p>HP offers a wide array of storage products and services to address the growing demands of the storage environment. HP Storage Services range from strategic design and planning, to integration and deployment, all the way to ongoing management capabilities. Storage Services can ensure the desired levels of performance, reliability, and availability, and they can ensure that storage meets the needs of the always-on Internet infrastructure.</p>
business recovery services	<p>For companies that depend on an always-on Internet infrastructure for their distributed IT environments, HP Business Recovery Services provide continuity or fast recovery of IT operations following natural and man-made disasters. These services complete the integrated mission-critical solution that helps ensure the continuity of a business in the face of all possible causes of unplanned downtime.</p>
outsourcing services	<p>HP's outsourcing services provide comprehensive, multivendor IT infrastructure management and application outsourcing services. These services are ideal for enterprises and service providers who face critical IT infrastructure challenges—and who want to use outsourcing to evolve their business models for a competitive advantage.</p>
<b>financing</b>	<p>HP's one-stop flexible financing services help customers accomplish their business objectives: improving cash flow, enhancing risk management, conserving capital, and speeding time to revenue.</p>
<b>for more information</b>	<p>For more information, please visit our Web site at <a href="http://www.hp.com/go/servers">www.hp.com/go/servers</a> and <a href="http://www.hp.com/go/rp7405">www.hp.com/go/rp7405</a></p>

Full information on HP products is available at [www.hp.com](http://www.hp.com)

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Windows is a U.S. registered trademark of Microsoft Corporation. Netscape is a U.S. trademark of Netscape Communications Corporation. Java is a U.S. trademark of Sun Microsystems Inc. UNIX is a registered trademark of The Open Group.

The information in this document is subject to change without notice.

© 2002 Hewlett-Packard Company

09/02

5981-3274EN

