


IBM System p5 560Q

Technical Overview and Introduction

Finer system granulation using Micro-Partitioning technology to help lower TCO

Modular midrange solution for managing on demand business

New POWER5+ processor options using DDR2 memory technology



Giuliano Anselmi
Bernard Filhol
SahngShin Kim
Gregor Linzmeier
Ondrej Plachy



International Technical Support Organization

**IBM System p5 560Q Technical Overview and
Introduction**

April 2007

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Archived

Second Edition (April 2007)

This edition applies to IBM System p5 560Q (product number 9116-561), Linux, and IBM AIX 5L Version 5.3, product number 5765-G03.

© Copyright International Business Machines Corporation 2006, 2007. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this Redpaper	ix
Become a published author	x
Comments welcome	x
Chapter 1. General description	1
1.1 System specifications	2
1.2 Physical package	3
1.3 Minimum and optional features	4
1.3.1 Processor card features	4
1.3.2 Memory features	5
1.3.3 Disk and media features	5
1.3.4 USB diskette drive	6
1.3.5 I/O drawers	6
1.3.6 Hardware Management Console models	8
1.4 Express Product Offerings	9
1.4.1 Express Product Offerings overview	10
1.4.2 Configurator starting points for Express Product Offerings	10
1.5 System racks	11
1.5.1 IBM 7014 Model T00 rack	11
1.5.2 IBM 7014 Model T42 rack	12
1.5.3 IBM 7014 Model S11 rack	13
1.5.4 IBM 7014 Model S25 rack	13
1.5.5 S11 rack and S25 rack considerations	13
1.5.6 The ac power distribution unit and rack content	14
1.5.7 Rack-mounting rules for p5-560Q	16
1.5.8 Useful Rack additions	17
1.5.9 OEM rack	19
Chapter 2. Architecture and technical overview	23
2.1 The POWER5+ processor	24
2.2 Processor and cache	25
2.2.1 p5-560Q POWER5+ Quad-Core Module	25
2.2.2 Available processor speeds	26
2.3 Memory subsystem	27
2.3.1 Memory placement rules	27
2.3.2 OEM memory	28
2.3.3 Memory throughput	29
2.4 System buses	29
2.4.1 Processor drawer interconnect cables	29
2.4.2 SP bus	30
2.4.3 RIO-2 buses, GX+ buses, and optional host adapters	30
2.5 Internal I/O subsystem	31
2.5.1 PCI-X slots and adapters	31
2.5.2 LAN adapters	32
2.5.3 SCSI adapters	33

2.5.4 Internal RAID options	33
2.5.5 iSCSI	33
2.5.6 Fibre Channel adapter	35
2.5.7 Graphic accelerator	36
2.5.8 InfiniBand Host Channel adapters	36
2.5.9 Asynchronous PCI-X adapters	36
2.5.10 Additional support for existing PCI-X adapters	36
2.5.11 Internal system ports	37
2.5.12 Ethernet ports	37
2.6 Internal storage	37
2.6.1 Internal hot-swappable SCSI disks	37
2.6.2 Internal media devices	38
2.7 External disk subsystems	38
2.7.1 IBM TotalStorage EXP24 Expandable Storage	38
2.7.2 IBM System Storage N3000 and N5000	39
2.7.3 IBM TotalStorage Storage DS4000 Series	39
2.7.4 IBM TotalStorage Enterprise Storage Server	39
2.8 Logical partitioning	40
2.8.1 Dynamic logical partitioning	40
2.9 Virtualization	40
2.9.1 POWER Hypervisor	41
2.10 Advanced POWER Virtualization feature	43
2.10.1 Micro-Partitioning technology	43
2.10.2 Logical, virtual, and physical processor mapping	44
2.10.3 Virtual I/O Server	46
2.10.4 Partition Load Manager	49
2.10.5 Integrated Virtualization Manager	49
2.11 Hardware Management Console	52
2.11.1 High availability using the HMC	53
2.11.2 IBM System Planning Tool	54
2.12 Operating system support	55
2.12.1 AIX 5L	55
2.12.2 Linux	57
2.13 Service information	58
2.13.1 Touch point colors	58
2.13.2 Operator control panel	58
2.13.3 System firmware	60
2.13.4 Service processor	63
2.13.5 Hardware management user interfaces	63
Chapter 3. RAS and manageability	67
3.1 Reliability, availability, and serviceability	68
3.1.1 Fault avoidance	68
3.1.2 First Failure Data Capture	68
3.1.3 Permanent monitoring	69
3.1.4 Self-healing	70
3.1.5 N+1 redundancy	71
3.1.6 Fault masking	71
3.1.7 Resource deallocation	71
3.1.8 Serviceability	72
3.2 Manageability	73
3.2.1 Service processor	73
3.2.2 Partition diagnostics	74

3.2.3 Service Agent	75
3.2.4 IBM System p5 firmware maintenance	77
3.3 Cluster solution	78
Related publications	81
IBM Redbooks	81
Other publications	81
Online resources	82
How to get IBM Redbooks	83
Help from IBM	83

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.


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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®

eServer™

pSeries®

AIX 5L™

AIX®

Chipkill™

DS4000™

DS6000™

DS8000™

Enterprise Storage Server®

FICON®

HACMP™

IntelliStation®

IBM®

Micro-Partitioning™

OpenPower™

PowerPC®

POWER™

POWER Hypervisor™

POWER4™

POWER5™

POWER5+™

PTX®

Redbooks®

RS/6000®

Service Director™

System p™

System p5™

System Storage™

TotalStorage®

Virtualization Engine™

1350™

The following terms are trademarks of other companies:

Data ONTAP, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Flex, and Portable Document Format (PDF) are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Internet Explorer, Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This document is a comprehensive guide that covers the IBM® System p5™ 560Q UNIX® servers. It introduces major hardware offerings and discusses their prominent functions.

Professionals wishing to acquire a better understanding of IBM System p™ products should read this document. The intended audience includes:

- ▶ Clients
- ▶ Sales and marketing professionals
- ▶ Technical support professionals
- ▶ IBM Business Partners
- ▶ Independent software vendors

This document expands the current set of IBM System p documentation by providing a desktop reference that offers a detailed technical description of the p5-560Q system.

This publication does not replace the latest IBM System p marketing materials and tools. It is intended as an additional source of information that you can use, together with existing sources, to enhance your knowledge of IBM server and workstation solutions.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Austin Center.

Giuliano Anselmi is a certified pSeries® Presales Technical Support Specialist working in the Field Technical Sales Support group based in Rome, Italy. For seven years, he was an IBM eServer™ pSeries Systems Product Engineer, supporting Web Server Sales Organization in EMEA, IBM Sales, IBM Business Partners, Technical Support Organizations, and IBM Dublin eServer Manufacturing. Giuliano has worked for IBM for 14 years, devoting himself to RS/6000® and pSeries systems with his in-depth knowledge of the related hardware and solutions.

Bernard Filhol is a UNIX Server Customer Satisfaction Resolution Team Leader for NEE and SWE IOTs in Montpellier France. He has more than 25 years of experience in mainframes and five years of experience in pSeries Client Satisfaction. He holds a degree in Electronic from Montpellier University Institute of Technology. His areas of expertise include Mainframe Channel Subsystem, FICON® and pSeries RAS. He has written extensively on FICON.

SahngShin Kim is a sales specialist on the STG infra-solution sales team in Seoul, Korea. He was previously a sales specialist for IBM eServer pSeries (three years) and grid computing. SahngShin has worked for IBM for six years, devoting himself to RS/6000 and pSeries systems and STG server products and their architecture.

Gregor Linzmeier is an IBM Advisory IT Specialist for RS/6000 and pSeries workstation and entry servers as part of the Systems and Technology Group in Mainz, Germany. He supports IBM sales, Business Partners, and clients with pre-sales consultation and implementation of client/server environments. He has worked for more than 15 years as an infrastructure specialist for RT, RS/6000, and AIX® in large CATIA client/server projects.

Ondrej Plachy is an IT specialist in IBM Czech Republic who is responsible for project design, implementation, and support of large scale computer systems. He has 11 years of

experience in the UNIX field. He holds the Ing. academic degree in Computer Science from Czech Technical University (CVUT), Prague. He worked at Supercomputing Centre of Czech Technical University for four years and has worked at IBM for seven years as a part of the AIX 5L™ support team.

The project that produced this document was managed by:

Scott Vetter
IBM US

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

Larry Amy, Baba Arimilli, Ron Arroyo, Joergen Berg, Terry Brennan, Erin Burke, Mark Dewalt, Bob Foster, Ron Gonzalez, Iggy Haider, Dan Henderson, Harmony Helterhof, David A. Hepkin, Tenley Jackson, Hal Jennings, Carolyn Jones, Brian J King, Bill Mihaltse, Thoi Nguyen, Ken Rozendal, Craig Shempert, Dave Willoughby
IBM

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbooks® document dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners or clients.

Your efforts will help increase product acceptance and client satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- Use the online **Contact us** review form found at:

ibm.com/redbooks

- Send your comments in an e-mail to:

redbook@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

General description

The IBM System p5 560Q rack-mount server is designed for greater application flexibility, with innovative technology, to capitalize the e-business revolution at the entry- or midrange- level server environments.

Note: With Express Product Offerings, based on the configuration, clients might qualify for a processor activation at no additional charge. The number of processor, total memory, quantity or capacity of disk storage, and presence of media device are the features that determine whether you are entitled to a processor activation. Contact your sales representative regarding the Express Product Offering or volume offerings.

The p5-560Q rack-mount server is a symmetric multiprocessor (SMP) system, based on a building block architecture and Quad-Core Module (QCM) technology. The system features 4-core, 8-core, or 16-core, 64-bit, copper and silicon on insulator (SOI)-based POWER5+™ microprocessors running at 1.5 GHz or 1.8 GHz packaged inside processor cards, together with L3 cache and memory DIMM sockets. Any processor card contains a 72 MB Level 3 cache, while main memory starts from 2 GB and can be expanded up to 128 GB, based on the DDR2 memory DIMMs at the time of writing, for higher performance and exploitation of 64-bit addressing to meet the demands of enterprise computing, such as large database applications.

Dynamic LPAR on a 16-core p5-560Q allows up to sixteen dedicated partitions. In addition, the optional Advanced POWER™ Virtualization feature enables the support of up to 160 micro-partitions. Micro-Partitioning™ technology enables multiple partitions to share a physical processor. The extended POWER Hypervisor™ controls dispatching the physical processors to each of the partitions using Micro-Partitioning technology. The Advanced POWER Virtualization feature allows sharing of processor capacity, physical network adapters and enables the virtualization of SCSI storage.

Simultaneous multithreading is a standard feature of POWER5+ technology that allows two threads to be executed at the same time on a single processor. Simultaneous multithreading is selectable with dedicated processors or processors in a shared pool using Micro-Partitioning technology. In addition, the POWER5+ microprocessor provides 90 nm chip design technology, which is designed to boost performance by allowing the QCM design.

A p5-560Q building block is featured as 4U high rack-mounted system and included six hot-plug PCI-X slots with Enhanced Error Handling (EEH) and six disk drive bays. A fully configured p5-560Q system is made up of two building blocks and features up to 128 GB of internal memory, 12 PCI-X slots, and 3.6 TB of internal storage capacity (based on 300 GB disk drives). Each building block, or system drawer, has two embedded Ultra320 SCSI controllers, one 10/100/1000 Mbps integrated dual-port Ethernet controller, two USB 2.0 capable ports, two slim-line media bays for DVD-RAM or DVD-ROM, and two System Power Control Network (SPCN) ports. Each complete system (either one or two building blocks) has two system ports, and two HMC ports. RIO attached I/O Drawers are only supported on the p5-560Q with the 1.8 GHz processor.

Additional reliability and availability features include redundant hot-swappable cooling fans and redundant power supplies. Along with these components, the p5-560Q is designed to provide an extensive set of reliability, availability, and serviceability (RAS) features that include: fault isolation, recovery from errors without stopping the system, avoidance of recurring failures, and predictive failure analysis.

The p5-560Q is backed by a three-year limited warranty. Check with your IBM representative for particular warranty availability in your region.

1.1 System specifications

Table 1-1 lists the general system specifications of the p5-560Q system.

Table 1-1 p5-560Q specifications

Description	Range
Operating temperature	5 to 35 degrees C (41 to 95 F)
Relative humidity	8% to 80%
Maximum wet bulb	23 degrees C (73.4 F) (operating)
Noise level	6.2 to 7.1 bels (operating 4-core configuration)
Operating voltage	200 to 240 V ac 50/60 Hz
Maximum power consumption	1,300 watts (maximum)
Maximum power source loading	1.37 kVA (maximum configuration)
Maximum thermal output	4,437 Btu ^a /hr (maximum configuration)

a. British Thermal Unit (BTU)

1.2 Physical package

One p5-560Q drawer, or building block, is packaged in a 4U¹ rack-mounted enclosure, and it is available only in the rack-mounted form factor. Table 1-2 lists the major physical attributes that are found on the p5-560Q building block.

Table 1-2 Physical packaging of the p5-560Q

Dimension	One p5-560Q building block
Height	174.1 mm (6.85 in.)
Width	483 mm (19.0 in.)
Depth	790 mm (31.1 in.)
Weight	63.6 kg (140 lb.)

Using the p5-560Q building block architecture, an installed system can be made of one to two building blocks. To help ensure the installation and serviceability in non-IBM, industry-standard racks, review the vendor's installation planning information for any product-specific installation requirements. The p5-560Q drawer interconnect cables present an additional planning requirement.

Figure 1-1 shows some views of the p5-560Q




Figure 1-1 Views of a p5-560Q

¹ One Electronic Industries Association Unit (1U) is 44.45 mm (1.75 in.).

1.3 Minimum and optional features

The p5-560Q server is a modular design based on POWER5+ processors located in Quad-Core Modules (QCM). The server can be made of up to two building blocks and is available in 4-core, 8-core, and 16-core configurations, featuring:

- ▶ Up to four QCM processor cards using the POWER5+ chips, for a total of 16 processors. The minimum requirement is one 4-core 1.5 GHz processor card with all processors activated. A 1.8 GHz processor is also available.
- ▶ From 2 GB to 128 GB of total system memory capacity using 533 MHz DDR2 DIMM technology.
- ▶ Up to 12 SCSI disk drives, six per drawer, for an internal storage capacity of 3.6 TB using 300 GB drives.
- ▶ 12 PCI-X slots, six per drawer. All PCI-X slots support Enhanced Error Handling (EEH). One GX+ adapter shares the same space with one short PCI-X slot.

Up to four slim-line media bays in two-drawers can be configured for optional optical storage devices.

The p5-560Q building block includes the service processor (SP) and the following native ports:

- ▶ Two 10/100/1000 Ethernet ports
- ▶ Two system ports
- ▶ Two USB 2.0 ports (external USB diskette drive 1.44 (FC 2591) is available.)
- ▶ Two HMC ports
- ▶ Two remote I/O (RIO-2) ports
- ▶ Two SPCN ports

In addition, the p5-560Q building block features two internal Ultra320 SCSI controllers, redundant hot-swappable power supply and redundant hot-swappable cooling fans, and redundant processor power regulators (FC 7768).

The combined system (made of two p5-560Q building blocks) requires the proper processor interconnect cable and the system SP Flex™ cable (see 2.4.1, "Processor drawer interconnect cables" on page 29, and 2.4.2, "SP bus" on page 30).

Note: In a p5-560Q combined system made of two building blocks, only the two HMC ports and the two system ports in the primary building block are available to use.

The system supports 32-bit and 64-bit applications and requires a specific levels of AIX 5L and Linux® operating systems. See 2.12, "Operating system support" on page 55.

When a system is ordered, changes to most system configuration features are available using an MES field upgrade.

1.3.1 Processor card features

The p5-560Q accommodates up to four 4-core processor cards running at 1.5 GHz (FC 7781) or 1.8 GHz (FC 7295) with 64-bit, copper-based, POWER5+ microprocessors. All installed processors, or cores, have to be activated. Any 4-core processor card contains a QCM that is made of two 2-core chips with 2 x 1.9 MB of L2 cache and 2 x 36 MB of L3 cache. Each

processor card accommodates eight slots for memory DIMMs using 533 MHz DDR2 technology.

An initial order must have at least one processor card. An order with two p5-560Q building blocks must have four processor cards, thus the maximum configuration is 16-core, but the 12-core configuration is not allowed. The server is available as 4-core, 8-core, and 16-core processor configurations.

1.3.2 Memory features

The minimum memory requirement for the p5-560Q server is 2 GB, and the maximum capacity is 128 GB, using 533 MHz DDR2 technology. The memory DIMM granularity supported in the p5-560Q is two. Table 1-3 lists the available memory features. Note that an amount of memory is always used by the POWER Hypervisor, even when the machine is not partitioned. The System Planning Tool can be used to calculate the amount of available memory for an operating system based on machine configuration (see 2.11.2, “IBM System Planning Tool” on page 54).

Table 1-3 Memory feature codes

Feature code	Description
1930	1 GB (2 x 512 MB) DIMMs, 276-pin DDR2, 533 MHz SDRAM
1931	2 GB (2 x 1 GB) DIMMs, 276-pin DDR2, 533 MHz SDRAM
1932	4 GB (2 x 2 GB) DIMMs, 276-pin DDR2, 533 MHz SDRAM
1934	8 GB (2 x 4 GB) DIMMs, 276-pin DDR2, 533 MHz SDRAM

We recommend that each processor card has an equal amount of memory installed. Balancing memory across the installed processor cards provides distributed memory access that might improve performance.

1.3.3 Disk and media features

Each p5-560Q building block is equipped with one 6-pack hot-pluggable disk drive enclosure divided into two logical 3-pack disk drive enclosures. The maximum internal storage capacity for a combined system made of two building blocks is 3.6 TB (using 300 GB Ultra320 disk drives, FC 1973, available at the time of writing).

The minimum p5-560Q configuration allowed includes six hot-pluggable disk drive slots and one disk. Table 1-4 shows the disk drive feature codes that each bay can contain.

Table 1-4 Hot-swappable disk drive options

Feature code	Description
1968	73.4 GB ULTRA320 10 K rpm SCSI hot-swappable disk drive
1969	146.8 GB ULTRA320 10 K rpm SCSI hot-swappable disk drive
1971	73.4 GB ULTRA320 15 K rpm SCSI hot-swappable disk drive
1972	146.8 GB ULTRA320 15 K rpm SCSI hot-swappable disk drive
1973	300 GB ULTRA320 10 K rpm SCSI hot-swappable disk drive

Each server drawer also features two slim-line media device bays. Any combination of DVD-ROM and DVD-RAM drives can be installed in any of the two slim-line bays:

- ▶ DVD-RAM drive, FC 1993
- ▶ DVD-ROM drive, FC 1994

A logical partition running a supported release of Linux requires a DVD-ROM drive or DVD-RAM drive to provide a way to run the stand-alone diagnostics from the diagnostics CD media. Concurrent diagnostics as provided by the AIX 5L **diag** command is not available on the Linux operating system at the time of writing.

Devices sharing the same media backplane must be assigned as a group to a single dynamic LPAR on a partitioned system.

1.3.4 USB diskette drive

The externally attached USB diskette drive provides storage capacity up to 1.44 MB (FC 2591) on high-density (2HD) floppy disks and 720 KB on a double density floppy disk. It includes a 350 mm (13.7 inch) cable with standard USB connector. This super-slim-line and lightweight USB V2-attached diskette drive takes its power requirements from the USB port. The drive can be attached to the integrated USB ports, or to a USB adapter (FC 2738). A maximum of one USB diskette drive is supported per integrated controller or adapter. The same controller can share a USB mouse and keyboard.

1.3.5 I/O drawers

System configurations of the p5-560Q built with the 1.8 GHz QCM processor card (FC 7295) will support the attachment of both the 7311-D11 and the 7311-D20 Remote I/O drawers using the two Remote I/O (RIO-2) native ports built into the CEC enclosure. Each CEC enclosure can support a maximum of four I/O drawers in a single loop, for a maximum of eight I/O drawers on a 16-core system. The 7311-D11 and 7311-D20 can be mixed in the same I/O loop.

7311 Model D11 I/O drawer

The 7311 Model D11 I/O drawer features six long PCI-X slots. Only the blind-swap cassettes are supported:

- ▶ FC 7862, for full-sized PCI cards
- ▶ FC 7861, for short-sized PCI cards
- ▶ FC 7863, for double-wide PCI cards

Two 7311 Model D11 I/O drawers fit side-by-side in the 4U enclosure (FC 7311) mounted in a 19-inch rack, such as the IBM 7014-T00 or 7014-T42.

The 7311 Model D11 I/O drawer offers a modular growth path for the p5-560Q systems with increasing I/O requirements.

The I/O drawer has the following attributes:

- ▶ 4U rack-mount enclosure (FC 7311) that can hold one or two D11 drawers
- ▶ Six PCI-X slots: 3.3 V, keyed, 133 MHz blind-swap hot-plug
- ▶ Default redundant hot-plug power and cooling devices
- ▶ Two RIO-2 and two SPCN ports

7311 Model D11 I/O drawer physical package

Because the 7311 Model D11 I/O drawer must be mounted into the rack enclosure (FC 7311), these are the physical characteristics of one I/O drawer or two I/O drawers side-by-side:

- ▶ One 7311 Model D11 I/O drawer
 - Width: 223 mm (8.8 in.)
 - Depth: 711 mm (28.0 in.)
 - Height: 175 mm (6.9 in.)
 - Weight: 19.6 kg (43 lb.)
- ▶ Two I/O drawers in a 7311 rack-mounted enclosure have the following characteristics:
 - Width: 445 mm (17.5 in.)
 - Depth: 711 mm (28.0 in.)
 - Height: 175 mm (6.9 in.)
 - Weight: 39.1 kg (86 lb.)

7311 Model D20 I/O drawer

The 7311 Model D20 I/O drawer is a 4U full-size drawer, which must be mounted in a rack. It features seven hot-pluggable PCI-X slots and, optionally, up to 12 hot-swappable disks arranged in two 6-packs. Redundant concurrently maintainable power and cooling is an optional feature (FC 6268).

PCI-X and PCI cards are inserted into the slot from the top of the I/O drawer. The installed adapters are protected by plastic separators, which are designed to prevent grounding and damage when adding or removing adapters.

The drawer has the following attributes:

- ▶ 4U rack mount enclosure assembly
- ▶ Seven PCI-X slots: 3.3 V, keyed, 133 MHz hot-plug
- ▶ Two 6-pack hot-swappable SCSI devices
- ▶ Optional redundant hot-plug power
- ▶ Two RIO-2 and two SPCN ports

Note: The 7311 Model D20 I/O drawer initial order, or an existing 7311 Model D20 I/O drawer that is migrated from another pSeries system, must have the RIO-2 ports available (FC 6417).

7311 Model D20 I/O drawer physical package

The I/O drawer has the following physical characteristics:

- ▶ Width: 482 mm (19.0 in.)
- ▶ Depth: 610 mm (24.0 in.)
- ▶ Height: 178 mm (7.0 in.)
- ▶ Weight: 45.9 kg (101 lb.)

Figure 1-2 shows the different views of the 7311-D20 I/O drawer.




Figure 1-2 7311-D20 I/O drawer views

Note: The 7311 Model D10, and the 7311 Model D11, or the 7311 Model D20 I/O drawers are designed to be installed by an IBM service representative.

I/O drawers and usable PCI slots

Table 1-5 summarizes the maximum number of I/O drawers supported and the total number of PCI-X slots available when expansion consists of a single drawer type.

Table 1-5 Maximum number of I/O drawers supported and total number of PCI slots

p5-560Q drawer	Max number of I/O drawers	Total number of PCI-X slots	
		D11	D20
1 drawer	4	30	34
2 drawer	8	54	62

^a One slot is reserved for the Remote I/O expansion card.

1.3.6 Hardware Management Console models

The p5-560Q can operate either using an Hardware Management Console (HMC) or without. When HMC managed, the HMC is required as a dedicated workstation that allows you to configure and manage partitions. The HMC provides a set of functions to manage the p5-560Q system LPARs, dynamic LPAR operations, virtual features, inventory and microcode management, and remote power control functions. These functions also include the handling of the partition profiles that define the processor, memory, and I/O resources allocated to an individual partition.

Note: Non-HMC managed modes are:

- ▶ Full system partition mode (only one partition containing all system resources exists on the system)
- ▶ Using the Integrated Virtualization Manager (IVM), see 2.10.5, “Integrated Virtualization Manager” on page 49

See 2.11, “Hardware Management Console” on page 52 for detailed information about the HMC.

Table 1-6 lists the last model available at the time of writing for the desktop and rack-mounted HMC options for POWER5™ processor-based systems. Existing HMC models can be also used when running the correct level of software.

Table 1-6 Latest HMC models available

Type-model	Description
7310-C05	IBM 7310 Model C05 desktide Hardware Management Console
7310-CR4	IBM 7310 Model CR4 rack-mount Hardware Management Console

Systems require Ethernet connectivity between the HMC and one of the Ethernet ports of service processor. Ensure that sufficient Ethernet adapters are available to enable public and private networks if you need both.

The 7310 Model C05 is a desktide model with only one native 10/100/1000 Ethernet port but can be extended with other two additional dual-port 10/100/1000 Gb Ethernet adapters.

The 7310 Model CR4 is a 1U, 19-inch rack mountable drawer that has two native Ethernet ports and can be extended with one additional two-port 10/100/1000 Gb Ethernet adapter.

In HMC managed installations with very high demand for high availability, two HMCs are recommended. The p5-560Q service processor allows the connection of two HMCs, so there are no additional features needed for a p5-560Q to support dual HMC environment. The HMCs provide a locking mechanism so that only one HMC at a time have write access to service processor.

When an HMC is connected to the p5-560Q, the integrated system ports are disabled. If you need serial connections, for example, non-Ethernet HACMP™ heartbeat, you need to provide an additional asynchronous adapter (see Table 2-6 on page 36).

Note: It is not possible to connect POWER4™ and POWER5 or POWER5+ processor-based systems simultaneously to the same HMC, but it is possible to connect POWER5 and POWER5+ processor-based systems together to the same HMC.

1.4 Express Product Offerings

The Express Product Offerings provide a convenient way to order a predefined configuration designed to meet typical client requirements. Special reduced pricing is available when a system order satisfies specific configuration requirements for memory, disk drives, and processors.

1.4.1 Express Product Offerings overview

When you order an Express, the configurator offers a choice of starting points that can be added onto. Clients can configure systems with one or two processor cards and four or eight processor activations.

With the purchase of an Express, for each paid processor activation the client is entitled to one processor activation at no additional charge, if the following requirements are met:

- ▶ The system must have at least two disk drives of at least 73.4 GB each.
- ▶ There must be at least 2 GB of memory installed for each active processor.

If you order a p5-560Q server Express Product Offering as defined here, you might qualify for a processor activation at no extra charge. The number of processors, total memory, quantity and size of disk, and presence of a media device are the only features that determine if a client is entitled to a processor entitlement at no additional charge.

When you purchase an Express, you are entitled to a lower priced AIX 5L or Linux operating system license or can choose to purchase the system with no operating system. The lower priced AIX 5L or Linux operating system is processed using a feature number on AIX 5L and either Red Hat or SUSE Linux. You can choose either the lower priced AIX 5L or Linux subscription, but not both. If you choose AIX 5L for your lower priced operating system, you can also order Linux but will purchase your Linux subscription at full price versus the reduced price. The same is true if you choose a Linux subscription as your lower priced operating system.

Systems with a reduced price AIX 5L offering are the IBM System p5 Express, AIX 5L edition; and systems with a lower priced Linux operating system will be referred to as the IBM System p5 Express, OpenPower™ edition. In the case of Linux, only the first subscription purchased is lower priced. So, for example, additional licenses purchased for Red Hat to run in multiple partitions are at full price.

You can make changes to the standard features as needed and still qualify for processor entitlements at no additional charge and a reduced price AIX 5L or Linux operating system license.

If the system was initially ordered as an Express, the system can be expanded at a later time using Express Product Offering pricing, when additional processors and activations along with the required memory are ordered on the same hardware upgrade order. The upgraded p5-560Q configuration must satisfy the Express Product Offering requirements for disk drives, memory, and processors. However, if the selection of total memory or disk drives are smaller than the total defined as the minimums, it disqualifies the order as an Express.

1.4.2 Configurator starting points for Express Product Offerings

The complete order must include the features identified in the minimum configuration plus the feature enhancements listed here:

- ▶ p5-560Q 4-core 1.5 GHz POWER5+ System
 - Processor (1 x FC 7781)
 - 8 GB memory (4 x FC 1931)
 - 2 x 73 GB DASD (2 x FC 1968)
 - Processor activations (2 x FC 7662, plus 2 x FC 7398)
 - Media backplane (1 x FC 7869)
 - IDE DVD-ROM (1 x FC 1994)

- ▶ p5-560Q 4-core 1.8 GHz POWER5+ System
 - Processor (1 x FC 7295)
 - 8 GB memory (4 x FC 1931)
 - 2 x 73 GB DASD (2 x FC 1968)
 - Processor activations (2 x FC 7235, plus 2 x FC 7345)
 - Media backplane (1 x FC 7869)
 - IDE DVD-ROM (1 x FC 1994)
- ▶ p5-560Q 8-core 1.5 GHz POWER5+ System
 - Processor (2 x FC 7781)
 - 16 GB memory (8 x FC 1931)
 - 2 x 73 GB DASD (2 x FC 1968)
 - Processor activations (4 x FC 7662, plus 4 x FC 7398)
 - Media backplane (1 x FC 7869)
 - IDE DVD-ROM (1 x FC 1994)
- ▶ p5-560Q 8-core 1.8 GHz POWER5+ System
 - Processor (2 x FC 7295)
 - 16 GB memory (4 x FC 1931)
 - 2 x 73 GB DASD (2 x FC 1968)
 - Processor activations (4 x FC 7235, plus 4 x FC 7345)
 - Media backplane (1 x FC 7869)
 - IDE DVD-ROM (1 x FC 1994)

1.5 System racks

The IBM 7014 Model S11, S25, T00, and T42 are 19-inch racks for general use with IBM System p rack-mount servers. The racks provide increased capacity, greater flexibility, and improved floor space utilization.

If a System p5 server is to be installed in a non-IBM rack or cabinet, you must ensure that the rack conforms to the EIA² standard EIA-310-D (see 1.5.9, “OEM rack” on page 19).

Note: It is the client's responsibility to ensure that the installation of the drawer in the preferred rack or cabinet results in a configuration that is stable, serviceable, safe, and compatible with the drawer requirements for power, cooling, cable management, weight, and rail security.

1.5.1 IBM 7014 Model T00 rack

The 1.8-meter (71-in.) Model T00 is compatible with past and present IBM System p systems. The T00 rack has the following features:

- ▶ 36 EIA units (36U) of usable space.
- ▶ Optional removable side panels.
- ▶ Optional highly perforated front door.
- ▶ Optional side-to-side mounting hardware for joining multiple racks.
- ▶ Standard business black or optional white color in OEM format.
- ▶ Increased power distribution and weight capacity.

² Electronic Industries Alliance (EIA). Accredited by American National Standards Institute (ANSI), EIA provides a forum for industry to develop standards and publications throughout the electronics and high-tech industries.

- ▶ Optional reinforced (ruggedized) rack feature (FC 6080) provides added earthquake protection with modular rear brace, concrete floor bolt-down hardware, and bolt-in steel front filler panels.
- ▶ Support for both ac and dc configurations.
- ▶ The dc rack height is increased to 1926 mm (75.8 in.) if a power distribution panel is fixed to the top of the rack.
- ▶ Up to four power distribution units (PDUs) can be mounted in the PDU bays (see Figure 1-3 on page 15), but others can fit inside the rack. See 1.5.6, “The ac power distribution unit and rack content” on page 14.
- ▶ An optional rack status beacon (FC 4690). This beacon is designed to be placed on top of a rack and cabled to servers, such as a p5-560Q, and other components inside the rack. Servers can be programmed to illuminate the beacon in response to a detected problem or changes in system status.
- ▶ A rack status beacon junction box (FC 4693) should be used to connect multiple servers to the beacon. This feature provides six input connectors and one output connector for the rack. To connect the servers or other components to the junction box or the junction box to the rack, status beacon cables (FC 4691) are necessary. Multiple junction boxes can be linked together in a series using daisy chain cables (FC 4692).
- ▶ Weights:
 - T00 base empty rack: 244 kg (535 pounds)
 - T00 full rack: 816 kg (1795 pounds)

1.5.2 IBM 7014 Model T42 rack

The 2.0-meter (79.3-inch) Model T42 addresses the client requirement for a tall enclosure to house the maximum amount of equipment in the smallest possible floor space. The features that differ in the Model T42 rack from the Model T00 include:

- ▶ 42 EIA units (42U) of usable space (6U of additional space).
- ▶ The Model T42 supports ac only.
- ▶ Weights:
 - T42 base empty rack: 261 kg (575 pounds)
 - T42 full rack: 930 kg (2045 pounds)

Optional Rear Door Heat eXchanger (FC 6858)

Improved cooling from the Rear Door Heat eXchanger enables clients to more densely populate individual racks freeing valuable floor space without the need to purchase additional air conditioning units. The Rear Door Heat eXchanger features:

- ▶ Water-cooled heat exchanger door designed to dissipate heat generated from the back of computer systems before it enters the room.
- ▶ An easy-to-mount rear door design that attaches to client-supplied water, using industry standard fittings and couplings.
- ▶ Up to 15 KW (approximately 50,000 BTUs/hr) of heat removed from air exiting the back of a fully populated rack.
- ▶ One year, limited warranty.

Physical specifications:

- ▶ Approximate height: 1945.5 mm (76.6 in.)
- ▶ Approximate width: 635.8 mm (25.03 in.)

- ▶ Approximate depth: 141.0 mm (5.55 in.)
- ▶ Approximate weight: 31.9 kg (70.0 lb.)

Client responsibilities:

- ▶ Secondary water loop (to building chilled water)
- ▶ Pump solution (for secondary loop)
- ▶ Delivery solution (hoses and piping)
- ▶ Connections: standard 3/4-inch internal threads

1.5.3 IBM 7014 Model S11 rack

The Model S11 rack satisfies many light-duty requirements for organizing smaller rack-mount servers and expansion drawers. The 0.6-meter-high rack has a perforated, lockable front door; a heavy-duty caster set for easy mobility; a complete set of blank filler panels for a finished look; EIA unit markings on each corner to aid assembly; and a retractable stabilizer foot. The Model S11 rack has the following specifications:

- ▶ Width: 520 mm (20.5 in.) with side panels
- ▶ Depth: 874 mm (34.4 in.) with front door
- ▶ Height: 612 mm (24.0 in.)
- ▶ Weight: 37 kg (75.0 lb.)

The S11 rack has a maximum load limit of 16.5 kg (36.3 lb.) per EIA unit for a maximum loaded rack weight of 216 kg (475 lb.).

1.5.4 IBM 7014 Model S25 rack

The 1.3-meter-high Model S25 rack satisfies many light-duty requirements for organizing smaller rack-mount servers. Front and rear rack doors include locks and keys, helping keep your servers secure. Side panels are a standard feature, simplifying ordering and shipping. This 25U rack can be shipped configured and can accept server and expansion units up to 28-inches deep.

The front door is reversible so that it can be configured for either left or right opening. The rear door is split vertically in the middle and hinges on both the left and right sides. The S25 rack has the following specifications:

- ▶ Width: 605 mm (23.8 in.) with side panels
- ▶ Depth: 1001 mm (39.4 in.) with front door
- ▶ Height: 1344 mm (49.0 in.)
- ▶ Weight: 100.2 kg (221.0 lb.)

The S25 rack has a maximum load limit of 22.7 kg (50 lb.) per EIA unit for a maximum loaded rack weight of 667 kg (1470 lb.).

1.5.5 S11 rack and S25 rack considerations

The S11 and S25 racks do not have vertical mounting space that can accommodate FC 7188 PDUs. All PDUs required for application in these racks must be installed horizontally in the rear of the rack. Each horizontally mounted PDU occupies 1U of space in the rack, and therefore reduces the space available for mounting servers and other components.

FC 0469 Customer Specified Rack Placement provides the client the ability to specify the physical location of the system modules and attached expansion modules (drawers) in the racks. The client's input is collected and verified through the IBM configurator tool and generate a configuration report file (CFReport). The client's request is reviewed by IBM

configurator tool for safe handling, by checking the weight distribution within the rack. At the end of this process, the IBM Manufacturing Plant provides the final approval for the configuration. This information is then used by IBM Manufacturing to assemble the system components (drawers) in the rack according to the client's request.

The CFReport must be submitted to the following site:

<http://www.ibm.com/servers/eserver/power/csp>

Table 1-7 lists the machine types and model numbers supported in the S11 and S25 racks.

Table 1-7 Models supported in S11 and S25 racks

Machine type-model	Name	Supported in	
		7014-S11 rack	7014-S25 rack
7037-A50	IntelliStation® POWER 185	Y	Y
7031-D24	EXP24 Disk Enclosure	Y	Y
7311-D20	I/O Expansion Drawer	Y	Y
9110-510	IBM eServer p5 510	Y	Y
9111-520	IBM eServer p5 520	Y	Y
9113-550	IBM eServer p5 550	Y	Y
9115-505	IBM System p5 505	Y	Y
9123-710	IBM eServer OpenPower 710	Y	Y
9124-720	IBM eServer OpenPower 720	Y	Y
9110-51A	IBM System p5 510 and 510Q	Y	Y
9131-52A	IBM System p5 520 and 520Q	Y	Y
9133-55A	IBM System p5 550 and 550Q	Y	Y
9116-561	IBM System p5 560		Y
9910-P33	3000 VA UPS (2700 watt)	Y	Y
9910-P65	500 VA UPS (208-240 V)		Y
7315-CR3	Rack-mount HMC		Y
7310-CR3	Rack-mount HMC		Y
7310-CR4	Rack-mount HMC		Y
7026-P16	LAN attached async. RAN		Y
7316-TF3	Rack-mounted flat-panel console kit		Y

1.5.6 The ac power distribution unit and rack content

For rack models T00 and T42, 12-outlet PDUs (FC 9188 and FC 7188) are available. For rack models S11 and S25, FC 7188 is also available.

Four PDUs can be mounted vertically in the T00 and T42 racks. See Figure 1-3 on page 15 for placement of the four vertically mounted PDUs. In the rear of the rack, two additional PDUs can be installed horizontally in the T00 rack and three in the T42 rack. The four vertical mounting locations will be filled first in the T00 and T42 racks. Mounting PDUs horizontally

consumes 1U per PDU and reduces the space available for other racked components. When mounting PDUs horizontally, we recommend that you use fillers in the EIA units occupied by these PDUs to facilitate proper air-flow and ventilation in the rack.

The S11 and S25 racks support as many PDUs as there is available rack space.

For detailed power cord requirements and power cord feature codes, see the publication *IBM System p5, eServer p5 and i5, and OpenPower Planning*, SA38-0508. For an online copy, see the IBM Systems Hardware Information Center. You can find it at:

<http://publib.boulder.ibm.com/eserver/>

Note: Ensure that the appropriate power cord feature is configured to support the power being supplied.

The Base/Side Mount Universal PDU (FC 9188) and the optional, additional, Universal PDU (FC 7188) support a wide range of country requirements and electrical power specifications. The PDU receives power through a UTG0247 power line connector. Each PDU requires one PDU-to-wall power cord. Nine power cord features are available for different countries and applications by varying the PDU-to-wall power cord, which must be ordered separately. Each power cord provides the unique design characteristics for the specific power requirements. To match new power requirements and save previous investments, these power cords can be requested with an initial order of the rack or with a later upgrade of the rack features.

The PDU has 12 client-usable IEC 320-C13 outlets. There are six groups of two outlets fed by six circuit breakers. Each outlet is rated up to 10 amps, but each group of two outlets is fed from one 15 amp circuit breaker.

Note: Based on the power cord that is used, the PDU can supply from 4.8 kVA to 19.2 kVA. The total kilovolt ampere (kVA) of all the drawers plugged into the PDU must not exceed the power cord limitation.

The Universal PDUs are compatible with previous models.




Figure 1-3 PDU placement and PDU view

Note: Each p5-560Q or a system drawer to be mounted in the rack, requires two power cords, which are not included in the base order. For maximum availability it is highly recommended to connect power cords from same p5-560Q or system drawer to two separate PDUs. These PDUs being connected to two independent client power sources.

1.5.7 Rack-mounting rules for p5-560Q

The primary rules that should be followed when mounting the p5-560Q into a rack are:

- ▶ The p5-560Q is designed to be placed at any location in the rack. For rack stability, it is advisable to start filling a rack from the bottom.

For p5-560Q configurations with two building blocks, all drawers must be installed together in the same rack, in a contiguous space of 8U within the rack.

- ▶ Any remaining space in the rack can be used to install other systems or peripherals, provided that the maximum permissible weight of the rack is not exceeded and the installation rules for these devices are followed.
- ▶ Before placing a p5-560Q into the service position, it is essential that the rack manufacturer's safety instructions have been followed regarding rack stability.
- ▶ Special consideration must be taken to avoid a flange on the top of the rack to clear the front bezel.
- ▶ When a p5-560Q system is installed in an IBM 7014-T00 or 7014-T42 rack that has no front door, a Thin Profile Front Trim Kit must be ordered for the rack. The front bezel on a model 560Q drawer is too wide to be used with the previously announced rack trim kits (FC 6107 and FC 6081). The required trim kit for the 7014-T00 rack is FC 6246. The required trim kit for the 7014-T42 rack is FC 6247.
- ▶ The IBM 7014-T42 rack is constructed with a small flange at the bottom of EIA location 37. This requires special placement rules when a model 560Q system is installed near the top of a 7014-T42 rack to avoid interference with the front bezel. No system drawer can be installed in EIA positions 34, 35, or 36. A two-drawer system cannot be installed above position 29. (The position number refers to the bottom of the lowest drawer.)

Depending on current implementation and future enhancements of additional I/O drawers connected to the p5-560Q or single installed p5-560Q systems, Table 1-8 shows examples of the minimum and maximum configurations for p5-560Q systems and attached I/O drawers.

Table 1-8 Minimum and maximum configurations for p5-560Qs and 7311-D20s

Rack models	One building block p5-560Q	Two building blocks p5-560Q
7014-T00	9	4
7014-T42	10	5
7014-S11	N/A	N/A
7014-S25	6	3

1.5.8 Useful Rack additions

This section highlights some solutions available to provide a single point of management for environments composed of multiple System p5-560Q servers or other IBM System p systems.

IBM 7212 Model 103 IBM TotalStorage storage device enclosure

The IBM 7212 Model 103 is designed to provide efficient and convenient storage expansion capabilities for selected IBM System p servers. The IBM 7212 Model 103 is a 1U rack-mountable option to be installed in a standard 19-inch rack using an optional rack-mount hardware feature kit. The 7212 Model 103 has two bays that can accommodate any of the following storage drive features:

- ▶ Digital Data Storage (DDS) Gen 5 DAT72 Tape Drive provides a physical storage capacity of 36 GB (72 GB with 2:1 compression) per data cartridge.
- ▶ VXA-2 Tape Drive provides a media capacity of up to 80 GB (160 GB with 2:1 compression) physical data storage capacity per cartridge.
- ▶ VXA-320 Tape Drive provides a media capacity of up to 160 GB (320 GB with 2:1 compression) physical data storage capacity per cartridge.
- ▶ DVD-RAM 2 drive can read and write on 4.7 GB and 9.4 GB DVD-RAM media. The DVD-RAM 2 uses only bare media and is also read compatible with multisession CD, CD-RW, and 2.6 GB and 5.2 GB DVD-RAM media.
- ▶ SLR60 Tape Drive provides a media capacity of up to 37.5 GB media (75 GB with 2:1 compression) physical data storage capacity per cartridge.
- ▶ SLR100 Tape Drive provides a media capacity of up to 50 GB media (100 GB with 2:1 compression) physical data storage capacity per cartridge.
- ▶ Half-High LTO-2 Tape Drive provides a media capacity of up to 200 GB media (400 GB with 2:1 compression) physical data storage capacity per cartridge.

Flat panel display options

The IBM 7316-TF3 Flat Panel Console Kit can be installed in the system rack. This 1U console uses a 17-inch thin film transistor (TFT) LCD with a viewable area of 337.9 mm x 270.03 mm and a 1280 x 1024 pels³ resolution. The 7316-TF3 Flat Panel Console Kit has the following attributes:

- ▶ A 17-inch, flat screen TFT color monitor that occupies only 1U (1.75 inches) in a 19-inch standard rack.
- ▶ Ability to mount the IBM Travel Keyboard in the 7316-TF3 rack keyboard tray.
- ▶ Support for the new 1x8 LCM switch (FC 4280), the Netbay LCM2 (FC 4279) with access to and control of as many as 64 servers, and support of both USB and PS/2 server-side keyboard and mouse connections.
- ▶ IBM Travel Keyboard mounts in the rack keyboard tray (Integrated Track point and UltraNav).

IBM PS/2 Travel Keyboards are supported on the 7316-TF3 for use in configurations where only PS/2 keyboard ports are available.

The IBM 7316-TF3 Flat Panel Console Kit provides an option for the USB Travel Keyboards with UltraNav. The keyboard enables the 7316-TF3 to be connected to systems that do not have PS/2 keyboard ports. The USB Travel Keyboard can be directly attached to an available

³ Picture elements

integrated USB port or a supported USB adapter (2738) on System p5 servers or 7310-CR4 and 7315-CR3 HMCs.

The IBM 7316-TF3 flat-panel, rack-mounted console is now available with two console switch options, which let you inexpensively cable, monitor, and manage your rack servers: the new 1x8 LCM Console Switch (FC 4280), and the LCM2 console switch (FC 4279).

The 1x8 Console Switch is a cost-effective, densely-packed solution that helps you set up and control selected System p rack-mounted IBM servers. Key features include the following:

- ▶ Supports one local user with PS/2 keyboard, PS/2 mouse, and video connections
- ▶ Features an 8-port, CAT5 console switch for single-user local management
- ▶ Supports both USB and PS/2 server-side keyboard and mouse connections
- ▶ Occupies only 1U (1.75 in) in a 19-in standard rack

The 1x8 Console Switch can be mounted in one of the following racks: 7014-T00, 7014-T42, 7014-S11, 7014-S25.

The 1x8 Console Switch supports GXT135P (FC 1980 and FC 2849) graphics accelerators. The following cables are used to attach the IBM servers to the 1x8 Console Switch:

- ▶ IBM 3M Console Switch Cable (PS/2) (FC 4282)
- ▶ IBM 3M Console Switch Cable (USB) (FC 4281)

The 1x8 Console Switch supports the following monitors:

- ▶ 7316-TF3 rack console monitor
- ▶ pSeries TFT monitors (FC 3641, 3643, 3644, 3645)

Separately available switch cables convert KVM signals for CAT5 cabling for servers with USB and PS/2 ports. A minimum of one cable feature (FC 4281) or USB feature (FC 4282) is required to connect the IBM 1x8 Console Switch (FC 4280) to a supported server. The 3 meter cable FC 4281 has one HD15 connector for video, and one USB connector for keyboard and mouse. The 3 meter cable FC 4282 has one HD15 connector for video, one PS/2 connector for a keyboard, and one PS/2 connector for a mouse. It is used to connect the IBM 1x8 Console Switch to a supported server.

The 1x8 Console Switch is a 1U (1.75-in) rack-mountable LCM switch containing eight analog rack interface ports for connecting switches using CAT5 cable. The switch supports a maximum video resolution of 1280x1024.

The Console Switch allows for two levels of tiering and supports up to 64 servers at a single user location through switch tiering. The previous VGA switch (FC 4200), the LCM (FC 4202) and LCM2 (FC 4279) switches can be tiered with the 1x8 Console Switch.

Note: When the 1x8 Console Switch is tiered with the previous VGA switch (FC 4200) or LCM (FC 4202) switch, it must be at the top level of the tier. When the 1x8 Console Switch is tiered with the LCM2 (FC 4279) switch, it must be at the secondary level of the tier.

The IBM Local 2x8 Console Manager (LCM2) switch (FC 4279) provides users single-point access and control of up to 1024 servers. It supports connection to servers with either PS/2 or USB connections with installation of appropriate options. Maximum resolution is 1280 x 1024 at 75Hz. The LCM2 switch can be tiered; three levels of tiering are supported.

A minimum of one LCM feature (FC 4268) or USB feature (FC 4269) is required with an IBM Local 2x8 Console Manager (LCM2) switch (FC 4279). Each feature can support up to four systems. When connecting to a p5-560Q, FC 4269 provides connection to the POWER5 USB ports. Only the PS/2 keyboard is supported when attaching the 7316-TF3 to the LCM Switch.

When selecting the LCM Switch, consider the following information:

- ▶ The KVM Conversion Option (KCO) cable (FC 4268) is used with systems with PS/2 style keyboard, display, and mouse ports.
- ▶ The USB cable (FC 4269) is used with systems with USB keyboard or mouse ports.
- ▶ The switch offers four ports for server connections. Each port in the switch can connect a maximum of 16 systems:
 - One KCO cable (FC 4268) or USB cable (FC 4269) is required for every four systems supported on the switch.
 - A maximum of 16 KCO cables or USB cables per port can be used with the Netbay LCM Switch to connect up to 64 servers.

Note: A server microcode update might be required on installed systems for boot-time System Management Services (SMS) menu support of the USB keyboards. For microcode updates, see:

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

We recommend that you have the 7316-TF3 installed between EIA 20 and EIA 25 of the rack for ease of use. The 7316-TF3 or any other graphics monitor requires a POWER GXT135P graphics accelerator (FC 1980 and FC 2849) to be installed in the server, or some other graphics accelerator, if supported.

Hardware Management Console 7310 Model CR4

The 7310 Model CR4 Hardware Management Console (HMC) is a 1U, 19-inch rack-mountable drawer that is supported in the 7014 racks. For additional HMC specifications, see 2.11, “Hardware Management Console” on page 52.

1.5.9 OEM rack

The p5-560Q can be installed in a suitable OEM rack, provided that the rack conforms to the EIA-310-D standard for 19-inch racks. This standard is published by the Electrical Industries Alliance, and a summary of this standard is available in the publication *IBM System p5, eServer p5 and i5, and OpenPower Planning*, SA38-0508.

The key points mentioned in this documentation are as follows:

- ▶ The front rack opening must be 451 mm wide + 0.75 mm (17.75 in. + 0.03 in.), and the rail-mounting holes must be 465 mm + 0.8 mm (18.3 in. + 0.03 in.) apart on center (horizontal width between the vertical columns of holes on the two front-mounting flanges and on the two rear-mounting flanges). See Figure 1-4 on page 20 for a top view showing the specification dimensions.




Figure 1-4 Top view of non-IBM rack specification dimensions

- The vertical distance between the mounting holes must consist of sets of three holes spaced (from bottom to top) 15.9 mm (0.625 in.), 15.9 mm (0.625 in.), and 12.67 mm (0.5 in.) on center, making each three-hole set of vertical hole spacing 44.45 mm (1.75 in.) apart on center. Rail-mounting holes must be 7.1 mm + 0.1 mm (0.28 in. + 0.004 in.) in diameter. See Figure 1-5 and Figure 1-6 on page 21 for the top and bottom front specification dimensions.




Figure 1-5 Rack specification dimensions, top front view




Figure 1-6 Rack specification dimensions, bottom front view

- ▶ It might be necessary to supply additional hardware, such as fasteners, for use in some manufacturer's racks.
- ▶ The system rack or cabinet must be capable of supporting an average load of 15.9 kg (35 lb.) of product weight per EIA unit.
- ▶ The system rack or cabinet must be compatible with drawer mounting rails, including a secure and snug fit of the rail-mounting pins and screws into the rack or cabinet rail support hole.

Note: The OEM rack must only support ac-powered drawers. We strongly recommend that you use a power distribution unit (PDU) that meets the same specifications as the PDUs to supply rack power. Rack or cabinet power distribution devices must meet the drawer power requirements, as well as the requirements of any additional products that will be connected to the same power distribution device.

Archived

Architecture and technical overview

This chapter discusses the overall system architecture represented by Figure 2-1 with its major components. The bandwidths provided throughout are theoretical maximums used for reference. You should always obtain real-world performance measurements using production workloads.

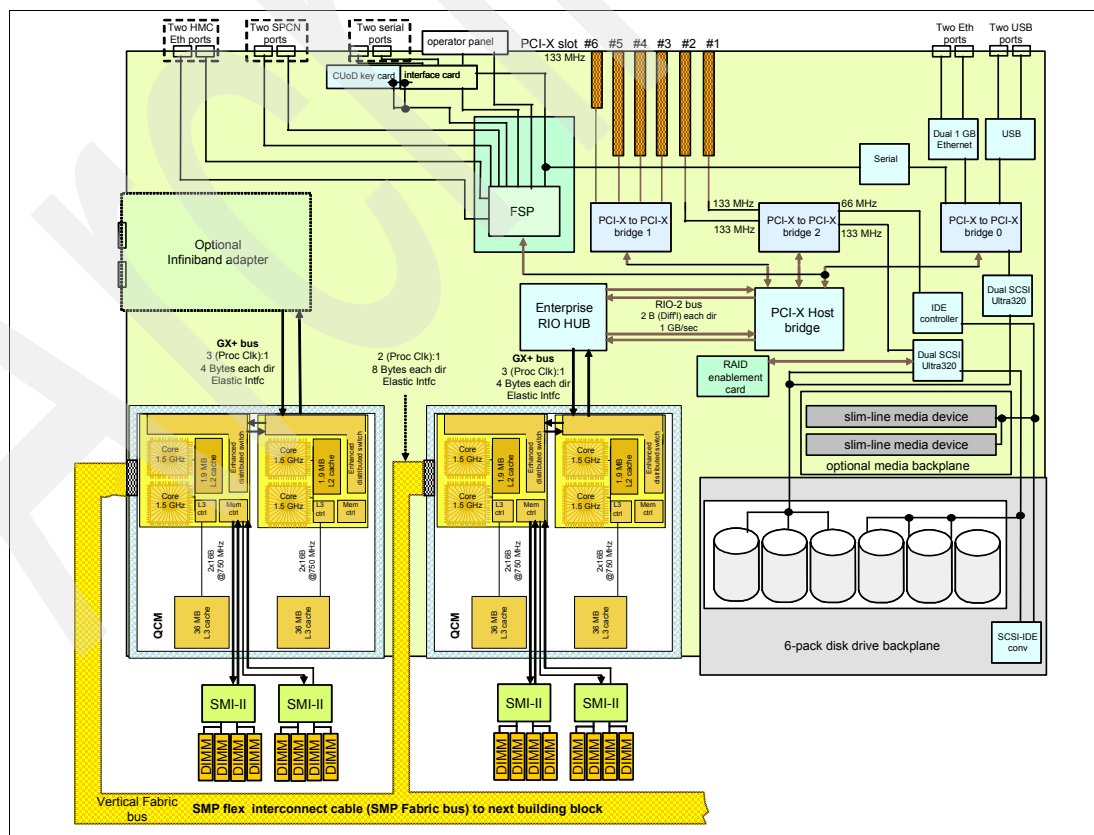


Figure 2-1 p5-560Q building block architecture shown with 1.5 GHz CPUs

2.1 The POWER5+ processor

The IBM POWER5+ chip capitalizes all the enhancements that are brought by the POWER5 chip. For a detailed description of the POWER5 chip, refer to *IBM System p5 570 Technical Overview and Introduction*, REDP-9117.

Figure 2-2 shows a high level view of the POWER5+ processor.

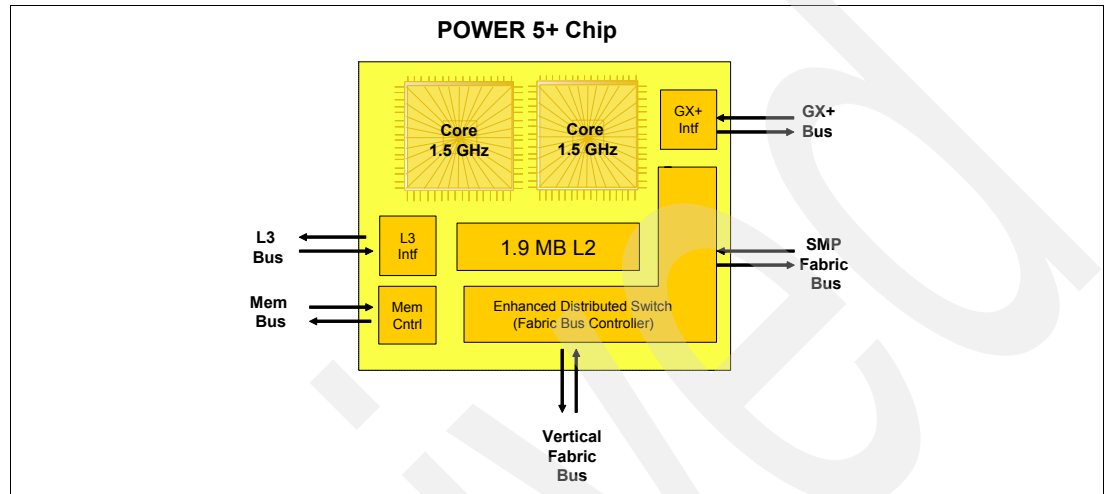


Figure 2-2 POWER5+ processor

The CMOS9S technology used in the POWER5 chip used a 130 nm fabrication process. The CMOS10S technology in the POWER5+ chip uses a 90 nm fabrication process, enabling:

- ▶ Performance gains through faster clock rates
- ▶ Chip size reduction (243 mm versus 389 mm)

Compared to the POWER5 chip, the 37 percent smaller POWER5+ chip consumes less power, thus requiring less sophisticated cooling. This allows its deployment in servers that previously only used low frequency chips due to cooling restrictions.

The POWER5+ design provides following additional enhancements over its predecessor:

- ▶ New pages sizes in ERAT and TLB. Two new pages sizes (64 KB and 16 GB) recently added in PowerPC® architecture.
- ▶ New segment size in SLB. One new segment size (1 TB) recently added in PowerPC architecture that is the common base architecture for POWER and PowerPC processors.
- ▶ The TLB size has been doubled in the POWER5+ over the POWER5 processor. The TLB in POWER5+ chip has 2048 entries.
- ▶ Floating-point round to integer instructions (frfin, frfiz, frfip, frfim) have been added to round floating-point numbers integers with the following rounding modes: nearest, zero, integer plus, integer minus.
- ▶ Improved floating-point performance.
- ▶ Lock performance enhancement.
- ▶ Enhanced SLB read.
- ▶ True Little-Endian mode. Support for the True Little-Endian mode as defined in the PowerPC architecture.

- ▶ Double the SMP support. Changes have been made in the fabric, L2 and L3 controller, memory controller, GX controller, and chip RAS to provide support for the QCM that allows the SMP system sizes to be double that available in POWER5 DCM-based servers. However current POWER5+ chip implementations only support a single address loop.
- ▶ Several enhancements have been made in the memory controller for improved performance. Ready to support DDR2 667 MHz DIMMs in the future.
- ▶ Enhanced redundancy in L1 Dcache, L2 cache and L3 directory. Independent control of the L2 cache and the L3 directory for redundancy to allow split-repair action has been added. More wordline redundancy has been added in the L1 Dcache. In addition, Array Built-In Self Test (ABIST) column repair for the L2 cache and the L3 directory has been added.

2.2 Processor and cache

In the p5-560Q, the POWER5+ chips, associated L3 cache chips, and memory DIMMs are packaged in a processor cards. The POWER5+ and L3 cache chips are packaged into the new POWER5+ QCM.

2.2.1 p5-560Q POWER5+ Quad-Core Module

The 4-core p5-560Q processor card contains a QCM and the local memory storage subsystem for that QCM. Two POWER5+ dual-core chips and their associated L3 cache chips are packaged in the QCM. Processor options include a 1.5 GHz or 1.8 GHz clock rate.

Note: Prior to installing a 1.8 GHz processor card, ensure the firmware is at level SF240-261 or later.

Figure 2-3 shows a layout view of p5-560Q QCM with associated memory.

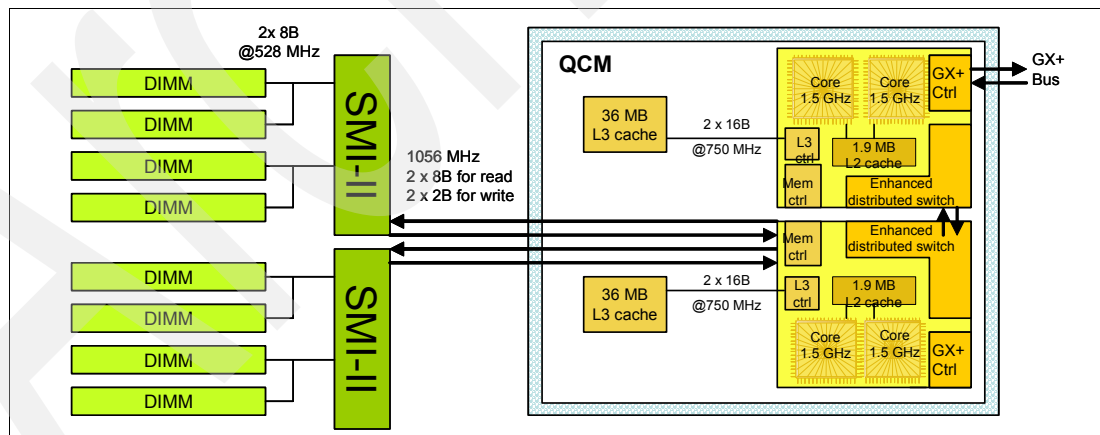


Figure 2-3 p5-560Q POWER5+ 1.5 GHz processor card with DDR2 memory socket layout view

The storage structure for the POWER5+ chip is a distributed memory architecture that provides high-memory bandwidth. Each processor in the QCM can address all memory and sees a single shared memory resource. In the QCM, only one POWER5+ chip has direct access to eight memory slots, controlled by two SMI-II chips, which are located in close physical proximity to the processor modules. The other POWER5+ chip has access to the same memory slots through the Vertical Fabric Bus.

The theoretical maximum throughput of the L3 cache is 16 byte read, 16 byte write at a bus frequency of 750 MHz (based on a 1.5 GHz processor clock) equates to 24000 MBps or 24 GBps.

I/O connects to the p5-560Q QCM using the GX+ bus. The QCM provides a single GX+ bus. Each processor in the POWER5+ chips has either a direct access to the GX+ bus using its GX+ bus controller or uses the Vertical Fabric Bus, controlled by the Fabric Bus controller. The GX+ bus provides an interface to I/O devices through the RIO-2 connections.

The POWER5+ chip that does not have direct access to memory does have a direct access to the GX+ bus. When more than one processor card is installed on a system (8-core or 16-core) the QCM to QCM communication is implemented using the Horizontal Fabric Bus.

Figure 2-4 shows a high level layout of a two processor card p5-560Q configuration.

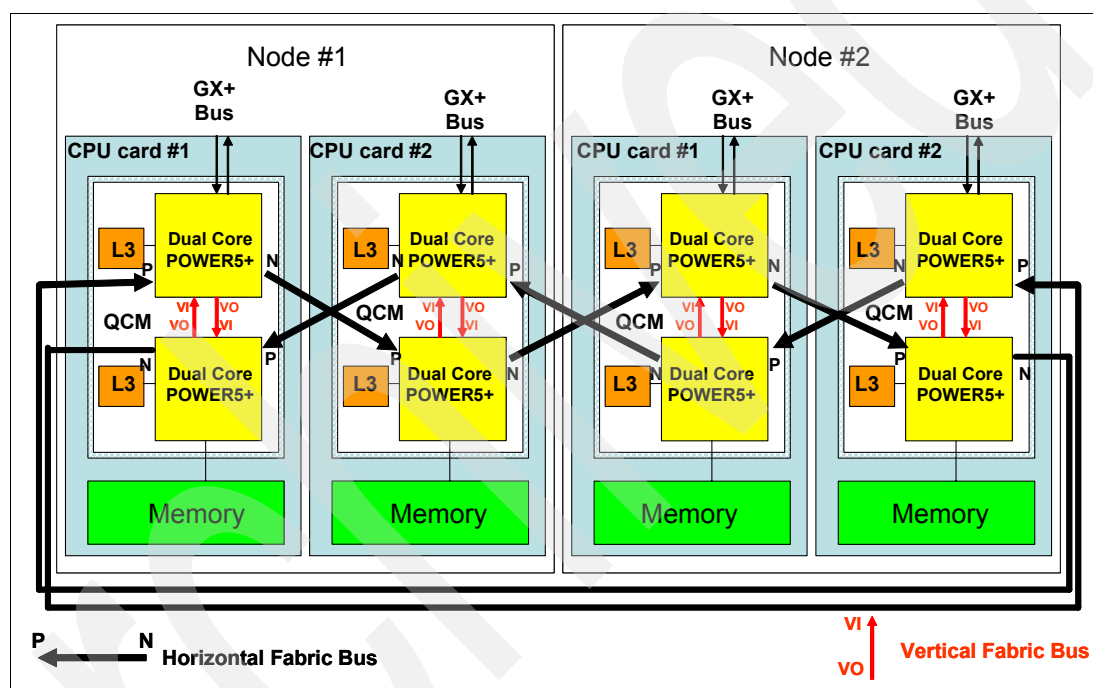


Figure 2-4 p5-560Q 16-core configuration

2.2.2 Available processor speeds

The 4-core to 16-core p5-560Q operates a processor clock rate of 1.5 GHz or 1.8 GHz.

To verify the processor characteristics on a running system, use one of the following commands:

► `lsattr -El procX`

Where *X* is the number of the processor, for example, `proc0` is the first processor in the system. The output from the command is similar to the following output (*False*, as used in this output, signifies that the value cannot be changed through an AIX 5L command interface):

frequency	1498500000	Processor Speed	False
smt_enabled	true	Processor SMT enabled	False
smt_threads	2	Processor SMT threads	False
state	enable	Processor state	False
type	powerPC_POWER5	Processor type	False

► **pmcycles -m**

The **pmcycles** command (available with AIX 5L) uses the performance monitor cycle counter and the processor real-time clock to measure the actual processor clock speed in MHz. The following output is from a 4-core p5-560Q system running at 1.5 GHz with simultaneous multithreading enabled:

```
Cpu 0 runs at 1498 MHz
Cpu 1 runs at 1498 MHz
Cpu 2 runs at 1498 MHz
Cpu 3 runs at 1498 MHz
Cpu 4 runs at 1498 MHz
Cpu 5 runs at 1498 MHz
Cpu 6 runs at 1498 MHz
Cpu 7 runs at 1498 MHz
```

Note: The **pmcycles** command is part of the **bos.pmapi** fileset. Use the **lspp -l bos.pmapi** command to determine if it is installed on your system.

2.3 Memory subsystem

Each processor card provides eight slots for up to eight pluggable DDR2 DIMMs. The DDR2 DIMMs has double rate bits comparing with DDR DIMMs (DDR DIMMs has double the data rate bits when compared with SDRAM) so that enables up to four times the performance of traditional SDRAM. Figure 2-5 on page 28 shows the memory slots and location codes. Every group of four memory slots are accessed by one Enhanced Synchronous Memory Interface (SMI-II) chip that is located between memory DIMMs and processor. The SMI-II supports multiple data flow modes.

2.3.1 Memory placement rules

Table 2-1 lists the memory features that are available for the p5-560Q at the time of writing.

Table 2-1 Available memory features

Feature code	Description
1930	1 GB (2 x 512 MB) DIMMs, 276-pin DDR2, 533 MHZ SDRAM
1931	2 GB (2 x 1 GB) DIMMs, 276-pin DDR2, 533 MHZ SDRAM
1932	4 GB (2 x 2 GB) DIMMs, 276-pin DDR2, 533 MHZ SDRAM
1934	8 GB (2 x 4 GB) DIMMs, 276-pin DDR2, 533 MHZ SDRAM

The memory can be plugged in pairs or quads, as required by the feature codes. Any memory feature available for a p5-560Q server consists of two DDR2 DIMMs. Different memory feature codes can be mixed within a system, but any quad must have the same memory DIMM. If the first memory pair is filled with the FC 1930, an additional memory feature must be another FC 1930 to make the quad. Mixed DIMM capacities in a quad will result in reduced RAS. The memory is installed in the first quad in the following order: J2A, J0A, J2C, and J0C; and for the second quad, in the order J2B, J0B, J2D, and J0D. In case more processor cards are present, memory must be balanced across the DIMM quad slots. The Service Information label, located on the top cover of the system, provides memory DIMMs slot location information.

To determine how much memory is installed in a system, use the following command:

```
# lsattr -El sys0 | grep realmem  
realmem      524288      Amount of usable physical memory in Kbytes      False
```

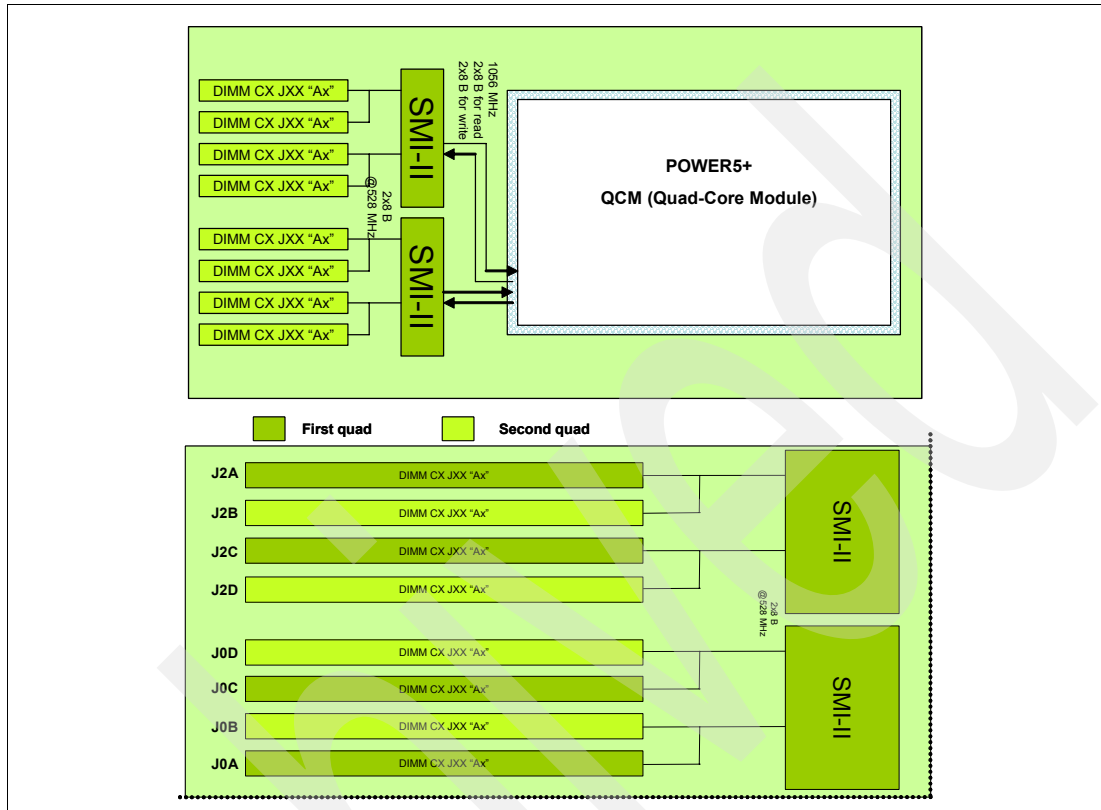


Figure 2-5 Memory placement for the p5-560Q server

2.3.2 OEM memory

OEM memory is not supported and not certified by IBM for use in a System p5 server. If the p5-560Q server is populated with OEM memory, you could experience unexpected and unpredictable behavior, especially when the system is using Micro-Partitioning technology.

All IBM memory is identified by an IBM logo and a white label printed with a barcode and an alphanumeric string, illustrated in Figure 2-6.




Figure 2-6 IBM memory certification label

2.3.3 Memory throughput

The memory subsystem throughput is depending on the speed of the memory. An elastic interface, contained in the POWER5+ chip, buffers reads and writes to and from memory and the processor. There are two SMI-II chips, each with a single 8-byte read and 2-byte write high speed elastic interface-II bus to the memory controller. A DDR2 bus allows double reads or writes per clock cycle. Because the bus operates at 1056 MHz, the peak processor-to-memory throughput for read is $(8 \times 2 \times 1056) = 16896$ MBps or 16.89 GBps. The peak processor-to-memory throughput for write is $(2 \times 2 \times 1056) = 4224$ MBps or 4.22 GBps, making total 21.12 GBps.

The DIMMs are 533 Hz DDR2 operating at 528 MHz through four 8-byte paths. Read and write operations share these paths. There must be at least four DIMMs installed to effectively use each path. In this case, the throughput between the SMI-II and the DIMMs is $(8 \times 4 \times 528)$ equal to 16.89 GBps.

2.4 System buses

This section provides additional information that is related to the internal buses.

2.4.1 Processor drawer interconnect cables

In systems made of two p5-560Q building blocks, the connection between processor cards in different building blocks is provided with a processor drawer interconnect cable.

Because of the redundancy and fault recovery built into the system interconnects, a drawer failure does not represent a system failure. When a problem is isolated and repaired, a system reboot might be required to re-establish full bus speed when the failure was specific to the interconnects.

The SMP fabric bus that connects the processors of separate p5-560Q building blocks is routed on the interconnect cable that is routed external to the building blocks. The flexible cable attaches directly to the processor cards, at the front of the p5-560Q building block, and is routed behind the front covers (bezels) of the p5-560Q building blocks. Figure 2-7 illustrates the logical fabric bus connections between the drawers.




Figure 2-7 Logical p5-560Q building blocks connection

2.4.2 SP bus

In addition to the processor drawer interconnect cable, see 2.4.1, “Processor drawer interconnect cables” on page 29, the interconnection of up to two p5-560Q building blocks requires the SP Flex cable (FC 1857) to ensure the vital data communications between the building blocks (see Figure 2-8). The SP Flex cable contains the system interconnect signals such as JTAG, I2C, clocks, and others.




Figure 2-8 SP Flex cables

2.4.3 RIO-2 buses, GX+ buses, and optional host adapters

Each QCM provides a GX+ bus that is used to connect to an I/O subsystem or Fabric Interface card. In a p5-560Q CEC drawer, there is one GX+ bus from each processor card slot. The second GX+ slot is not active unless the second processor card is installed. An optional hot-pluggable GX Dual-port 4x InfiniBand Host Channel Adapter (FC 1810) can be added to connect the Topspin Server Switches (see 2.5.8, “InfiniBand Host Channel adapters” on page 36 for more information). It must be installed starting with the first p5-560Q building block, and the related PCI-X slot 6 must remain empty.

The p5-560Q provides two external RIO-2 ports, which can operate up to 1 GHz. The RIO-2 ports are used for I/O expansion to external I/O drawers. The supported I/O drawers are the 7311 Model D11 and 7311 Model D20.

Table 2-2 provides a summary of the cumulative throughput rates for memory, cache, and I/O for the p5-560Q.

Table 2-2 Memory, processor cache, and I/O throughput

Processor speed (GHz)	Core	Memory (GB/s)	L2 to L3 (GB/s)	GX+ (GB/s)
1.5	4	21.1	48	4
1.5	8	42.2	96	8
1.5	16	84.4	192	16
1.8	4	21.1	57.8	4.8
1.8	8	42.2	115.2	9.6
1.8	16	84.4	230.4	19.2

2.5 Internal I/O subsystem

An internal RIO-2 bus is imbedded in the system planar. The system planar contains the Enterprise RIO-2 hub, the PCI-X Host bridge chip to connect the six hot-plug PCI-X slots and the integrated I/O, that is packaged on the system planar, through the PCI-X to PCI-X bridges. The connections of the PCI-X slots and integrated I/O devices to the PCI-X to PCI-X bridges are properly distributed to maximize the system performances.

All the PCI-X slots, except PCI-X slot 6, are full length 64-bit capable running at 133 MHz. PCI-X slot 6 can accept only short cards and share its space with the optional GX Dual-port 4x InfiniBand Host Channel Adapter (FC 1810).

2.5.1 PCI-X slots and adapters

PCI-X, where the X stands for extended, is an enhanced PCI bus, delivering a bandwidth of up to 1 Gbps, running a 64-bit bus at 133 MHz. PCI-X is backward compatible, so the p5-560Q systems can support existing 3.3 volt PCI adapters.

The PCI-X slots in the p5-560Q system support hot-plug and Extended Error Handling (EEH). In the unlikely event of a problem, EEH-enabled adapters respond to a special data packet generated from the affected PCI-X slot hardware by calling system firmware, which will examine the affected bus, allow the device driver to reset it, and continue without a system reboot. In the p5-560Q, all the PCI-X and PCI adapters are inside a carrier that provides a blind-swap mechanism to facilitate the addition, the replacement or the removal of any hot-pluggable PCI-X or PCI adapter.

Blind-swap carrier

Blind-swap carriers allow any PCI-X or PCI adapter to slide into the rear of a server. The carrier is designed so that the adapter is *guided* into place on a set of rails. The adapter is seated in the slot, completing the electrical connection by shifting an attached lever. This capability allows the PCI-X or PCI adapters to be concurrently replaced without having to pull out the server from the rack. and place it in a service position.

All PCI-X slots on a p5-560Q are PCI 2.2-compliant and are hot-plug enabled, which allows most PCI adapters to be removed, added, or replaced without powering down the system. This function enhances 24 x 7 system availability and serviceability when the operating system is configured to take advantage of dual paths through hardware. Figure 2-9 shows the blind-swap carrier that is used in the p5-560Q.




Figure 2-9 Blind-swap carrier

64-bit and 32-bit adapters

IBM offers 64-bit adapter options for the p5-560Q, as well as 32-bit adapters. Higher-speed adapters use 64-bit slots because they can transfer 64 bits of data for each data transfer phase. Generally, 32-bit adapters can function in 64-bit PCI-X slots. However, some 64-bit adapters cannot be used in 32-bit slots. For a full list of the adapters that are supported on the p5-560Q systems, and for important information regarding adapter placement, see the IBM Systems Hardware Information Center. You can find it at:

<http://publib.boulder.ibm.com/eserver/>

The p5-560Q internal PCI-X slots support a wide range of PCI-X I/O adapters to handle your I/O requirements.

2.5.2 LAN adapters

To connect a p5-560Q to a local area network (LAN), you can use the dual port internal 10/100/1000 Mbps RJ-45 Ethernet controller, integrated on the system planar.

Table 2-3 lists additional LAN adapters that are available for an initial system order at the time of writing. IBM supports an installation with NIM using Ethernet and token-ring adapters (CHRP¹ is the platform type).

Table 2-3 Available LAN adapters

Feature code	Adapter description	Type	Slot	Size	Max ^a
1954	4-port 10/100/1000 Ethernet	Copper	32 or 64	short	6
1978	Gigabit Ethernet	Fiber	32 or 64	short	6
1979	Gigabit Ethernet	Copper	32 or 64	short	6
1983	2-port Gigabit Ethernet	Copper	32 or 64	short	6
1984	2-port Gigabit Ethernet	Fiber	32 or 64	short	6
5721	10 Gigabit Ethernet-SR PCI-X	Fiber	64	short	

¹ CHRP stands for Common Hardware Reference Platform, a specification for PowerPC-based systems that can run multiple operating systems.

Feature code	Adapter description	Type	Slot	Size	Max ^a
5722	10 Gigabit Ethernet-LR PCI-X	Fiber	64	short	

a. The column reports the maximum number of adapters per p5-560Q building block

2.5.3 SCSI adapters

To connect to external SCSI devices, the adapters provided in Table 2-4 are available, at the time of writing, to be configured with a p5-560Q system initial order.

Table 2-4 Available SCSI adapters

Feature code	Adapter description	Slot	Size	Max ^a
1912	Dual Channel Ultra320 SCSI	64	short	5
1913	Dual Channel Ultra320 SCSI RAID	64	long	5

a. The column reports the maximum number of adapters per p5-560Q building block

Note: At the time of writing, previous SCSI adapters are also supported to be used in the p5-560Q but cannot be part of an initial order configuration. Clients that would like to connect existing external SCSI devices can contact the IBM service representative

For more information about the internal SCSI system, see 2.6, “Internal storage” on page 37.

2.5.4 Internal RAID options

Every p5-560Q building block system is delivered with a disk drive cage that supports up to six disk drive units, offering both internal RAID and non-RAID solutions in any building block. When internal RAID solution is not required, at least one 73.4 GB 10 K disk drive (FC 1968), or 36.4 GB 15 K disk drive (FC 1970) is required.

The internal RAID solution requires at least three disk drives and the SCSI RAID Enablement Card (FC 5728). When the SCSI RAID Enablement Card is installed in the system, it re-sequences the two SCSI controllers that support the six disk drive bays, transforming the system from two logical 3-packs of disk drives to one physical 6-pack of disk drives. It does this by disabling one of the integrated SCSI controllers. This also adds the requirement that all six disks be assigned to a single LPAR when the card is present.

The RAID implementation requires a minimum of three disk drives to form a RAID array, so when an order comes in place with FC 5728, at least three disk drives must be in the order list.

Note: Because the p5-560Q building block has six disk drive bays, clients performing upgrades must plan accordingly to ensure the correct handling of their RAID arrays.

2.5.5 iSCSI

iSCSI is an open, standards-based approach by which SCSI information is encapsulated using the TCP/IP protocol to allow its transport over IP networks. It allows transfer of data between storage and servers in block I/O formats (defined by iSCSI protocol) and thus enables the creation of IP SANs. iSCSI allows an existing network to transfer SCSI commands and data with full location independence and defines the rules and processes to accomplish the communication. The iSCSI protocol is defined in iSCSI IETF draft-20.

For more information about this standard, see:

<http://tools.ietf.org/html/rfc3720>

Although iSCSI can be, by design, supported over any physical media that supports TCP/IP as a transport, today's implementations are only on Gigabit Ethernet. At the physical and link level layers, iSCSI supports Gigabit Ethernet and its frames so that systems supporting iSCSI can be directly connected to standard Gigabit Ethernet switches and IP routers. iSCSI also enables the access to block-level storage that resides on Fibre Channel SANs over an IP network using iSCSI-to-Fibre Channel gateways such as storage routers and switches.

The iSCSI protocol is implemented on top of the physical and data-link layers and presents to the operating system standard SCSI Access Method command set. It supports SCSI-3 commands and reliable delivery over IP networks. The iSCSI protocol runs on the host initiator and the receiving target device. It can either be optimized in hardware for better performance on an iSCSI host bus adapter (such as FC 1986 and FC 1987 supported in IBM System p5 servers) or run in software over a standard Gigabit Ethernet network interface card. IBM in System p5 systems support iSCSI in the following two modes:

Hardware	Using iSCSI adapters (see "IBM iSCSI adapters" on page 34)
Software	Supported on standard Gigabit adapters, additional software (see "IBM iSCSI software Host Support Kit" on page 35) must be installed. The main processor is use for processing related to iSCSI protocol

Initial iSCSI implementations are targeted at small to medium-sized businesses and departments or branch offices of larger enterprises that have not deployed Fibre Channel SANs. iSCSI is an affordable way to create IP SANs from a number of local or remote storage devices. If there is Fibre Channel present, as it is typically in a data center, it can be accessed by the iSCSI SANs (and vice versa) via iSCSI-to-Fibre Channel storage routers and switches.

iSCSI solutions always involve the following software and hardware components:

Initiators	These are the device drivers and adapters that reside on the client. They encapsulate SCSI commands and route them over the IP network to the target device.
Targets	The target software receives the encapsulated SCSI commands over the IP network. The software can also provide configuration support and storage-management support. The underlying target hardware can be a storage appliance that contains embedded storage, it can also be a gateway or bridge product that contains no internal storage of its own.

IBM iSCSI adapters

New iSCSI adapters in IBM System p5 systems provide advantage of the increased bandwidth through the hardware support of iSCSI protocol. The 1 Gigabit iSCSI TOE PCI-X adapters support hardware encapsulation of SCSI commands and data into TCP and transports over the Ethernet using IP packets. The adapter operates as an iSCSI TOE (TCP/IP Offload Engine). This offload function eliminates host protocol processing and reduces CPU interrupts. Adapter uses Small form factor LC type fiber optic connector or copper RJ45 connector.

Table 2-5 provides the orderable iSCSI adapters.

Table 2-5 Available iSCSI adapters

Feature code	Description	Slot	Size	Max
1986	Gigabit iSCSI TOE PCI-X on copper media adapter	64	short	3

Feature code	Description	Slot	Size	Max
1987	Gigabit iSCSI TOE PCI-X on optical media adapter	64	short	3

IBM iSCSI software Host Support Kit

The iSCSI protocol can also be used over standard Gigabit Ethernet adapters. To use this approach, download the appropriate iSCSI Host Support Kit for your operating system from the IBM NAS support Web site at:

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

The iSCSI Host Support Kit on AIX 5L and Linux acts as software iSCSI initiator and allows to access iSCSI target storage devices using standard Gigabit Ethernet network adapters. To ensure the best performance enable the TCP Large Send, TCP send and receive flow control, and Jumbo Frame features of the Gigabit Ethernet Adapter and the iSCSI Target. Tune network options and interface parameters for maximum iSCSI I/O throughput on the operating system.

IBM System Storage N series

The combination of System p5 and IBM System Storage™ N Series as the first of a whole new generation of iSCSI enabled storage products provide an end-to-end set of solutions. Currently the System Storage N series feature three models: N3700, N5200, and N5500.

All models provide:

- ▶ Support for entry-level and midrange customers requiring Network Attached Storage (NAS) or Internet Small Computer System Interface (iSCSI) functionality.
- ▶ Support for Network File System (NFS), Common Internet File System (CIFS), and iSCSI protocols
- ▶ Data ONTAP® software (at no charge), with plenty of additional functions such as data movement, consistent snapshots and NDMP server protocol, some available through optional licensed functions.
- ▶ Enhanced reliability with optional clustered (2-node) failover support.

2.5.6 Fibre Channel adapter

The p5-560Q supports up to six 2 Gigabit Fibre Channel PCI-X Adapters (FC 1977) in any building block. The PCI-X adapter is a 64-bit, short form factor adapter with an LC type external fibre connector that provides single or dual initiator capability over an optical fiber link or loop. With the use of appropriate optical fiber cabling, this adapter provides the capability for a network of high speed local and remote located storage. Distances of up to 500 meters running at 1 Gbps data rate and up to 300 meters running at 2 Gbps data rate are supported between the adapter and an attaching device or switch. When used with IBM supported Fibre Channel storage switches supporting long-wave optics, distances of up to 10 kilometers are capable running at either 1 Gbps or 2 Gbps data rates.

The 2 Gigabit Fibre Channel PCI-X Adapter can be used to attach devices either directly, or using the supported Fibre Channel Switches. If attaching a device or switch with an SC type fibre connector, also the LC-SC 50 Micron Fiber Converter Cable (FC 2456) or a LC-SC 62.5 Micron Fiber Converter Cable (FC 2459) is required.

2.5.7 Graphic accelerator

The p5-560Q supports up to eight enhanced POWER GXT135P (FC 1980) 2D graphic accelerators. The POWER GXT135P is a low-priced 2D graphics accelerator for IBM System p5 servers. This adapter supports both analog and digital monitors.

The available graphic accelerator option is supported for System Management Services (SMS), firmware, and other functions, as well as when AIX 5L or Linux starts an X11-based graphic user interface (GUI).

2.5.8 InfiniBand Host Channel adapters

The p5-560Q supports one GX Dual-port 4x InfiniBand Host Channel Adapter (FC 1810), per building block, that enables the attachment of the Topspin Server Switch models 120 and 270. The GX Dual-port 4x InfiniBand HCA plugs into the system planar, using the GX slot. Connection to the Topspin Server Switches are accomplished by using the 4x IB Cables. If the GX Dual-port 4x InfiniBand HCA is present, the PCI-X slot 6 cannot be used.

Topspin Server Switch models 120 and 270

Switches are the fundamental components of an InfiniBand fabric. An IBM System p5 server proposal might also include the Topspin Server Switch model 120 and 270 in an initial system order.

The Topspin 120 and 270 Server Switch are a programmable switching platform that consists of a switched multiterabit interconnect and an intelligent control architecture. The high-bandwidth, low-latency interconnection is extremely adaptable. The switches enable an outstanding level of application scaling, rapid deployment, and resource consolidation.

See the following link for more Topspin Server Switch information:

<http://www.topspin.com/solutions/index.htm>

2.5.9 Asynchronous PCI-X adapters

The asynchronous PCI-X adapters provide connection of asynchronous EIA-232 or RS-422 devices. The system ports are not supported for serial HACMP heartbeat. Table 2-6 provides a list of the available adapters that support serial HACMP heartbeat:

Table 2-6 Asynchronous PCI-X adapters

Feature code	Description
2943	8-Port Asynchronous Adapter EIA-232/RS-422
5723	2-Port Asynchronous IEA-232 PCI Adapter (9-pin ^a)

^a In many cases, the 5723 async adapter is configured to supply a backup HACMP heartbeat. In these cases, a FC 3124 or FC 3125 serial cable must be also configured. Since the adapter provides a 9-pin connector and the cable provides a 25-pin connector, the FC 3925 converter is required to attach the cable to the adapter.

2.5.10 Additional support for existing PCI-X adapters

The lists of the major PCI-X adapters that can be configured in a p5-560Q when an initial configuration order is going to be built are described in 2.5.2, "LAN adapters" on page 32 to 2.5.9, "Asynchronous PCI-X adapters" on page 36. The list of all the supported PCI-X adapters, with the related support for additional external devices, is more extended.

Clients that would like to use existing PCI-X adapters can contact the IBM service representative to verify if supported.

2.5.11 Internal system ports

The system ports S1 and S2, at the rear of the system, are only available if the system is not managed using a Hardware Management Console (HMC). In this case, the S1 and S2 ports provide limited support of serial consoles or modems.

If an HMC is connected, a *virtual serial console* is provided by the HMC (logical device vsa0 under AIX) and also a modem can be connected to the HMC. The S1 and S2 ports are not usable in this case.

If serial port function is needed, optional PCI adapters are available, see 2.5.9, “Asynchronous PCI-X adapters” on page 36.

2.5.12 Ethernet ports

The two built-in Ethernet ports provide 10/100/1000 Mbps connectivity over CAT-5 cable for up to 100 meters. Table 2-7 lists the attributes of the LEDs that are visible on the side of the jack.

Table 2-7 Ethernet LED descriptions

LED	Light	Description
Link	Off Green	No link; could indicate a bad cable, not selected, or configuration error Connection established
Activity	On Off	Data activity Idle

2.6 Internal storage

Two Ultra320 SCSI controllers under EADS-X chips that are integrated into the system planar per building block are used to drive the internal disk drives. The six internal drives plug into the disk drive enclosure, which has two separate SCSI controllers with three disk drives each. Each of these controllers can be dynamically assigned to partitions if required.

The internal disk drive bays can be used in two different modes, depending on whether the SCSI RAID Enablement Card (FC 5728) is installed (See 2.5.4, “Internal RAID options” on page 33.).

2.6.1 Internal hot-swappable SCSI disks

Up to six hot-swappable disk drives plug in the two logical 3-pack disk drive backplanes of a p5-560Q building block. The hot-swap process is controlled by the virtual SCSI Enclosure Services (VSES), which is located in the logical 3-pack disk drive backplane. (AIX 5L assigns the name vses0 to the first 3-pack, and vses1 to the second). The two logical 3-pack disk drive logical enclosures can accommodate the devices listed in Table 2-8.

Table 2-8 Available disk drives

Feature code	Description
1971	73.4 GB ULTRA320 15 K rpm SCSI hot-swappable disk drive

Feature code	Description
1972	146.8 GB ULTRA320 15 K rpm SCSI hot-swappable disk drive
1973	300 GB ULTRA320 10 K rpm SCSI hot-swappable disk drive

Prior to the hot-swap of a disk drive in the hot-swappable-capable bay, all necessary operating system actions must be undertaken to ensure that the disk is capable of being deconfigured. After the disk drive has been deconfigured, the SCSI enclosure device will power-off the slot, enabling safe removal of the disk. You should ensure that the appropriate planning has been given to any operating-system-related disk layout, such as the AIX 5L Logical Volume Manager, when using disk hot-swap capabilities. For more information, see *Problem Solving and Troubleshooting in AIX 5L*, SG24-5496.

Note: It is recommended that you follow this procedure, after the disk has been deconfigured, when removing a hot-swappable disk drive:

1. Release the tray handle on the disk assembly.
2. Pull out the disk assembly a little bit from the original position.
3. Wait up to 20 seconds until the internal disk stops spinning.

Now, you can safely remove the disk from the DASD backplane.

After the SCSI disk hot-swap procedure, you can expect to find `SCSI_ERR10` logged in the AIX 5L error log, with the second word of the sense data equal to 0017. It is generated from a SCSI bus reset that is issued by the VSES to reset all processes when a drive is inserted, and it is not an issue. These errors can be removed using the `errclear -N SCSI-ERR10 0` command.

Hot-swappable disks and Linux

Hot-swappable disk drives on IBM System p5 systems are supported with SUSE Linux Enterprise Server 9 for POWER, or later, and Red Hat Enterprise Linux AS for POWER Version 3, or later.

2.6.2 Internal media devices

The p5-560Q provides two slim-line media bays per drawer (building block) for optional DVD-ROM (FC 1994) and optional DVD-RAM (FC 1993). These devices share the same controller and therefore must be assigned together to the same LPAR.

2.7 External disk subsystems

The p5-560Q has internal hot-swappable drives. When the AIX 5L operating system is installed in a IBM System p5 server, the internal disks are usually used for the AIX 5L rootvg volume group and paging space. Specific client requirements can be satisfied with the several external disk possibilities that the p5-560Q supports.

2.7.1 IBM TotalStorage EXP24 Expandable Storage

The IBM TotalStorage® EXP24 Expandable Storage disk enclosure, Model D24 or T24, can be purchased together with the p5-560Q and will provide low-cost Ultra320 (LVD) SCSI disk storage. This disk storage enclosure device provides more than 7 TB of disk storage in a 4U rack-mount (Model D24) or compact deskside (Model T24) unit. Whether high availability

storage solutions or simply high capacity storage for a single server installation, the unit provides a cost-effective solution. It provides 24 hot-swappable disk bays, 12 accessible from the front and 12 from the rear. Disk options that can be accommodate in any of the four six-packs disk drive enclosure are 73.4 GB, 146.8 GB, or 300 GB 10 K rpm or 36.4 GB, 73.4 GB, or 146.8 GB 15 K rpm drives. Each of the four six-packs disk drive enclosure might be attached independently to an Ultra320 SCSI or Ultra320 SCSI RAID adapter. For high available configurations, a dual bus repeater card (FC 5742) allows each six-pack to be attached to two SCSI adapters, installed in one or multiple servers or logical partitions. Optionally, the two front or two rear six-packs might be connected together to form a single Ultra320 SCSI bus of 12 drives.

2.7.2 IBM System Storage N3000 and N5000

The IBM System Storage N3000 and N5000 line of iSCSI enabled storage offerings provide a flexible way to implement a Storage Area Network over an Ethernet network. Flexible-Fibre Channel and SATA disk drive capabilities allow for deployment in multiple solution environments including data compliant retention, nearline storage, disk-to-disk backup scenarios and high-performance mission-critical I/O intensive operations. See the following link for more information:

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

2.7.3 IBM TotalStorage Storage DS4000 Series

The IBM System Storage DS4000™ line of Fibre Channel enabled Storage offerings provides a wide range of storage solutions for your Storage Area Network. The IBM TotalStorage DS4000 Storage server family consists of the following models: DS4100, DS4300, DS4500, and DS4800. The Model DS4100 Express Model is the smallest model and scales up to 44.8 TB; the Model DS4800 is the largest and scales up to 89.6 TB of disk storage at the time of this writing. Model DS4300 provides up to 16 bootable partitions, or 64 bootable partitions if the turbo option is selected, that are attached with the Gigabit Fibre Channel Adapter (FC 1977). Model DS4500 provides up to 64 bootable partitions. Model DS4800 provides 4 GB switched interfaces. In most cases, both the IBM TotalStorage DS4000 family and the IBM System p5 servers are connected to a storage area network (SAN). If only space for the rootvg is needed, the Model DS4100 is a good solution.

For support of additional features and for further information about the IBM TotalStorage DS4000 Storage Server family, refer to the following Web site:

<http://www.ibm.com/servers/storage/disk/ds4000/index.html>

2.7.4 IBM TotalStorage Enterprise Storage Server

The IBM TotalStorage Enterprise Storage Server® (ESS) Models DS6000™ and DS8000™ are the high-end premier storage solution for use in storage area networks and use POWER technology-based design to provide fast and efficient serving of data. The IBM TotalStorage DS6000 provides enterprise class capabilities in a space-efficient modular package. It scales to 67.2 TB of physical storage capacity by adding storage expansion enclosures. The Model DS8000 series is the flagship of the IBM TotalStorage DS family. The DS8000 scales to 192 TB. However, the system architecture is designed to scale to over one petabyte. The Model DS6000 and DS8000 systems can also be used to provide disk space for booting LPARs or partitions using Micro-Partitioning technology. ESS and the IBM System p5 servers are usually connected together to a storage area network.

For further information about ESS, refer to the following Web site:

http://www.ibm.com/servers/storage/disk/enterprise/ds_family.html

2.8 Logical partitioning

Dynamic logical partitions (LPARs) and virtualization increase utilization of system resources and add a new level of configuration possibilities. This section provides details and configuration specifications about this topic. The virtualization discussion includes virtualization enabling technologies that are standard on the system, such as the POWER Hypervisor, and optional ones, such as the Advanced POWER Virtualization feature.

2.8.1 Dynamic logical partitioning

Logical partitioning (LPAR) was introduced with the POWER4 processor-based product line and the AIX 5L Version 5.1 operating system. This technology offered the capability to divide a pSeries system into separate logical systems, allowing each LPAR to run an operating environment on dedicated attached devices, such as processors, memory, and I/O components.

Later, dynamic LPAR increased the flexibility, allowing selected system resources, such as processors, memory, and I/O components, to be added and deleted from dedicated partitions while they are executing. AIX 5L Version 5.2, with all the necessary enhancements to enable dynamic LPAR, was introduced in 2002. The ability to reconfigure dynamic LPARs encourages system administrators to dynamically redefine all available system resources to reach the optimum capacity for each defined dynamic LPAR.

Operating system support for dynamic LPAR

Table 2-9 lists AIX 5L and Linux support for dynamic LPAR capabilities.

Table 2-9 Operating system supported function

Function	AIX 5L Version 5.2	AIX 5L Version 5.3	Linux SLES 9	Linux RHEL AS 3	Linux RHEL AS 4
Dynamic LPAR capabilities (add, remove, and move operations)					
Processor	Y	Y	Y	N	Y
Memory	Y	Y	N	N	N
I/O slot	Y	Y	Y	N	Y

2.9 Virtualization

With the introduction of the POWER5 processor, partitioning technology moved from a dedicated resource allocation model to a virtualized shared resource model. This section briefly discusses the key components of virtualization on System p servers.

For more information about virtualization, see the following Web site:

<http://www.ibm.com/servers/eserver/about/virtualization/systems/pseries.html>

You can also consult the following IBM Redbooks:

<http://www.redbooks.ibm.com/abstracts/sg247940.html?open>

<http://www.redbooks.ibm.com/abstracts/sg245768.html?open>

2.9.1 POWER Hypervisor

Combined with features designed into the POWER5 and POWER5+ processors, the POWER Hypervisor delivers functions that enable other system technologies, including Micro-Partitioning technology, virtualized processors, IEEE VLAN, compatible virtual switch, virtual SCSI adapters, and virtual consoles. The POWER Hypervisor is a basic component of system firmware that is always active, regardless of the system configuration.

The POWER Hypervisor provides the following functions:

- ▶ Provides an abstraction between the physical hardware resources and the logical partitions that uses them.
- ▶ Enforces partition integrity by providing a security layer between logical partitions.
- ▶ Controls the dispatch of virtual processors to physical processors (see later discussion in 2.10.2, “Logical, virtual, and physical processor mapping” on page 44).
- ▶ Saves and restores all processor state information during logical processor context switch.
- ▶ Controls hardware I/O interrupt management facilities for logical partitions.
- ▶ Provides virtual LAN channels between physical partitions that help to reduce the need for physical Ethernet adapters for inter-partition communication.

The POWER Hypervisor is always active when the server is running partitioned or not and also when not connected to the HMC. It requires memory to support the logical partitions on the server. The amount of memory required by the POWER Hypervisor firmware varies according to several factors. Factors influencing the POWER Hypervisor memory requirements include the following:

- ▶ Number of logical partitions
- ▶ Partition environments of the logical partitions
- ▶ Number of physical and virtual I/O devices used by the logical partitions
- ▶ Maximum memory values given to the logical partitions

Note: Use the LPAR Validation Tool for estimate the memory requirements of the POWER Hypervisor.

In AIX 5L V5.3, the **lparstat** command using the **-h** and **-H** flags displays the POWER Hypervisor statistical data. Using the **-h** flag adds summary POWER Hypervisor statistics to the default **lparstat** output.

The minimum amount of physical memory for each partition is 128 MB, but in most cases the actual requirements and recommendations are between 256 MB and 512 MB for AIX 5L, Red Hat and Novell SUSE Linux. Physical memory is assigned to partitions in increments of Logical Memory Block (LMB). For POWER5+ processor-based systems, LMB might be adjusted from 16 MB to 256 MB.

The POWER Hypervisor provides the following types of virtual I/O adapters:

- ▶ Virtual SCSI
- ▶ Virtual Ethernet
- ▶ Virtual (TTY) console

Virtual SCSI

The POWER Hypervisor provides virtual SCSI mechanism for virtualization of storage devices (a special logical partition to install the Virtual I/O Server is required to use this

feature, as described in 2.10.3, “Virtual I/O Server” on page 46). The storage virtualization is accomplished using two, paired, adapters: a virtual SCSI server adapter and a virtual SCSI client adapter. Only the Virtual I/O Server partition can define virtual SCSI server adapters, other partitions are *client* partitions. The Virtual I/O Server is available with the optional Advanced POWER Virtualization feature (FC 7304).

Virtual Ethernet

The POWER Hypervisor provides a virtual Ethernet switch function that allows partitions on the same server to use a fast and secure communication without any need for physical interconnection. The virtual Ethernet allows a transmission speed in the range of 1 to 3 GBps, depending on the MTU² size and CPU entitlement. Virtual Ethernet requires system with either AIX 5L Version 5.3 or appropriate level of Linux supporting virtual Ethernet devices (see chapter 2.12, “Operating system support” on page 55). The virtual Ethernet is part of the base system configuration.

Virtual Ethernet has the following major features:

- ▶ The virtual Ethernet adapters can be used for both IPv4 and IPv6 communication and can transmit packets with a size up to 65408 bytes. Therefore, the maximum MTU for the corresponding interface can be up to 65394 (65390 if VLAN tagging is used).
- ▶ The POWER Hypervisor presents itself to partitions as a virtual 802.1Q compliant switch. Maximum number of VLANs is 4096. virtual Ethernet adapters can be configured as either untagged or tagged (following IEEE 802.1Q VLAN standard).
- ▶ A partition supports 256 virtual Ethernet adapters. Besides a default port VLAN ID, the number of additional VLAN ID values that can be assigned per Virtual Ethernet adapter is 20, which implies that each Virtual Ethernet adapter can be used to access 21 virtual networks.
- ▶ Each partition operating system detects the virtual local area network (VLAN) switch as an Ethernet adapter without the physical link properties and asynchronous data transmit operations.

Any virtual Ethernet can also have connection outside of the server if a layer-2 bridging to a physical Ethernet adapter is set in one Virtual I/O server partition (see 2.10.3, “Virtual I/O Server” on page 46 for more details about shared Ethernet).

Note: Virtual Ethernet is based on the IEEE 802.1Q VLAN standard. No physical I/O adapter is required when creating a VLAN connection between partitions, and no access to an outside network is required.

Virtual (TTY) console

Each partition needs to have access to a system console. Tasks such as operating system installation, network setup, and some problem analysis activities require a dedicated system console. The POWER Hypervisor provides the virtual console using a virtual TTY or serial adapter and a set of Hypervisor calls to operate on them. Virtual TTY does not require the purchase of any additional features or software such as the Advanced POWER Virtualization feature.

Depending on the system configuration, the operating system console can be provided by the Hardware Management Console virtual TTY, IVM virtual TTY, or from a terminal emulator that is connected to a system port.

² Maximum transmission unit

2.10 Advanced POWER Virtualization feature

The Advanced POWER Virtualization feature (FC 7304) is an optional, additional cost feature. This feature enables the implementation of more fine-grained virtual partitions on IBM System p5 servers.

The Advanced POWER Virtualization feature includes:

- ▶ Firmware enablement for Micro-Partitioning technology.
Support for up to 10 partitions per processor using 1/100 of the processor granularity. Minimum CPU requirement per partition is 1/10. All processors are enabled for micro-partitions (number of processors on system equals the number of Advanced POWER Virtualization features ordered).
- ▶ Installation image for the Virtual I/O Server software that is shipped as a system image on DVD. Client partitions can be either AIX 5.3 or Linux. It supports:
 - Ethernet adapter sharing (Ethernet bridge from virtual Ethernet to external network).
 - Virtual SCSI Server.
 - Partition management using Integrated Virtualization Manager (Virtual I/O Server Version 1.2 or later only).
- ▶ Partition Load Manager (AIX 5L Version 5.3 only)
 - Automated CPU and memory reconfiguration.
 - Real-time partition configuration and load statistics.
 - Graphical user interface.

For more details about Advanced POWER Virtualization and virtualization in general, see:

<http://www.ibm.com/servers/eserver/pseries/ondemand/ve/resources.html>

2.10.1 Micro-Partitioning technology

The concept of Micro-Partitioning technology allows you to allocate fractions of processors to the partition. The Micro-Partitioning technology is only available with POWER5 and POWER5+ processor-based systems. From an operating system perspective, a virtual processor cannot be distinguished from a physical processor, unless the operating system has been enhanced to be made aware of the difference. Physical processors are abstracted into virtual processors that are available to partitions. See 2.10.2, “Logical, virtual, and physical processor mapping” on page 44 for more details.

When defining a shared partition, several options have to be defined:

- ▶ Minimum, desired, and maximum processing units. Processing units are defined as processing power, or fraction of time, the partition is dispatched on physical processors.
- ▶ The processing sharing mode, either capped or uncapped.
- ▶ Weight (preference) in the case of uncapped partition.
- ▶ Minimum, desired, and maximum number of virtual processors.

The POWER Hypervisor calculates partition's processing *entitlement* based on minimum, desired, and maximum values, sharing mode and also based on other active partitions' requirements. The actual entitlement is never smaller than the minimum value but can exceed the maximum value in case of uncapped partition.

A partition can be defined with a processor capacity as small as 0.10 processing units. This represents one-tenth of a physical processor. Each physical processor can be shared by up to 10 shared processor partitions and partition's entitlement can be incremented fractionally by as little as one-hundredth of the processor. The shared processor partitions are dispatched and time-sliced on the physical processors under control of the POWER Hypervisor. The shared processor partitions are created and managed by the HMC or Integrated Virtualization Management (included with Virtual I/O Server software version 1.2 or later). There is only one pool of shared processors at the time of writing this publication and all shared partitions are dispatched by the Hypervisor within this pool. Dedicated partitions and micro-partitions can coexist on the same POWER5+ processor-based server as long as enough processors are available.

The p5-560Q supports up to a 16-core configuration, therefore up to sixteen dedicated partitions, or up to 160 micro-partitions can be created. It is important to point out that the maximums stated are supported by the hardware, but the practical limits depend from the application workload demands.

2.10.2 Logical, virtual, and physical processor mapping

The meaning of the term *physical processor* in this section is a *processor core*. For example, in a 2-core server with a DCM (Dual-Core Module) there are two physical processors, in a 4-core configuration with a QCM (Quad-Core Module) there are four physical processors.

In dedicated mode, physical processors are assigned as a whole to partitions. The simultaneous multithreading feature in the POWER5+ processor core allows the core to execute instructions from two independent software threads simultaneously. To support this feature, the concept of *logical processors* was introduced. Operating system (AIX 5L or Linux) sees one physical processor as two logical processors if the simultaneous multithreading feature is on. It can be turned off while operating system is executing (for AIX 5L, use the `smtctl` command). If simultaneous multithreading is off then each physical processor is presented as one logical processor and thus only one thread is executed on the physical processor at the time.

In a micro-partitioned environment with shared mode partitions, an additional concept of *virtual processors* was introduced. Shared partitions can define any number of virtual processors (maximum number is 10 times the number of processing units assigned to the partition). From the POWER Hypervisor point of view, the virtual processors represent dispatching objects (for example, the POWER Hypervisor dispatches virtual processors to physical processors according to partition's processing units entitlement). At the end of the POWER Hypervisor's dispatch cycle (10 ms), all partitions should receive total CPU time equal to their processing units entitlement. Virtual processors are either running (dispatched) on a physical processor or standby (waiting). An operating system is able to dispatch its software threads to these virtual processors and is completely screened from actual number of physical processors. The logical processors are defined on top of virtual processors in the same way as though they are physical processors. So, even with a virtual processor, the concept of logical processor exists and the number of logical processor depends whether the simultaneous multithreading is turned on or off.

Some additional information related to the virtual processors is as follows:

- ▶ There is one-to-one mapping of running virtual processors to physical processors at any given time. The number of virtual processors that can be active at any given time cannot exceed the total number of physical processors in shared processor pool.
- ▶ A virtual processor can be either running (dispatched) on a physical processor or standby waiting for a physical processor to become available

- ▶ Virtual processors do not introduce any additional abstraction level, they are really only a dispatch entity. When running on a physical processor, virtual processors run at the same speed as the physical processor.
- ▶ Each partition's profile defines CPU entitlement that determines how much processing power any given partition should receive. Total sum of CPU entitlement of all partitions cannot exceed number of available physical processors in shared processor pool.
- ▶ A partition has an amount of processing power regardless of the number of virtual processors that it defines.
- ▶ A partition can use more processing power, regardless of its entitlement, if it is defined as an *uncapped* partition in the partition profile. If there is spare processing power available in shared processor pool or other partitions are not using their entitlement, an uncapped partition can use additional processing units if its entitlement is not enough to satisfy its application processing demand in the given processing entitlement.
- ▶ When the partition is uncapped, the number of defined virtual processors determines the limitation of the maximum processing power it can receive. For example if number of virtual processors is two, then the maximum usable processor units is two.
- ▶ It is allowed to define more virtual processors than physical processors. In that case, virtual processor will be waiting for dispatch more often and some performance impact caused by redispersing virtual processors on physical processors should be considered. It is also true that some applications can benefit from using more virtual processors than the physical processors.
- ▶ The number of virtual processors can be changed dynamically through a dynamic LPAR operation.

Virtual processor recommendations

For each partition you can define a number of virtual processors set to the maximum processing power the partition could ever request. If there are, for example, four physical processors installed in the system, one production partition and three test partitions, then:

- ▶ Define production LPAR with four virtual processors so that it can receive full processing power of all four physical processors during the time the other partitions are idle.
- ▶ If you know that the test system will never consume more than one processor computing unit, then they should be defined with one virtual processor. Some test systems might require additional virtual processors, such as four, in order to use idle processing power left over by a production system during off-business hours.

Figure 2-10 shows logical, virtual, and physical processor mapping, and an example of how the virtual processor and logical processor can be dispatched to the physical processor.

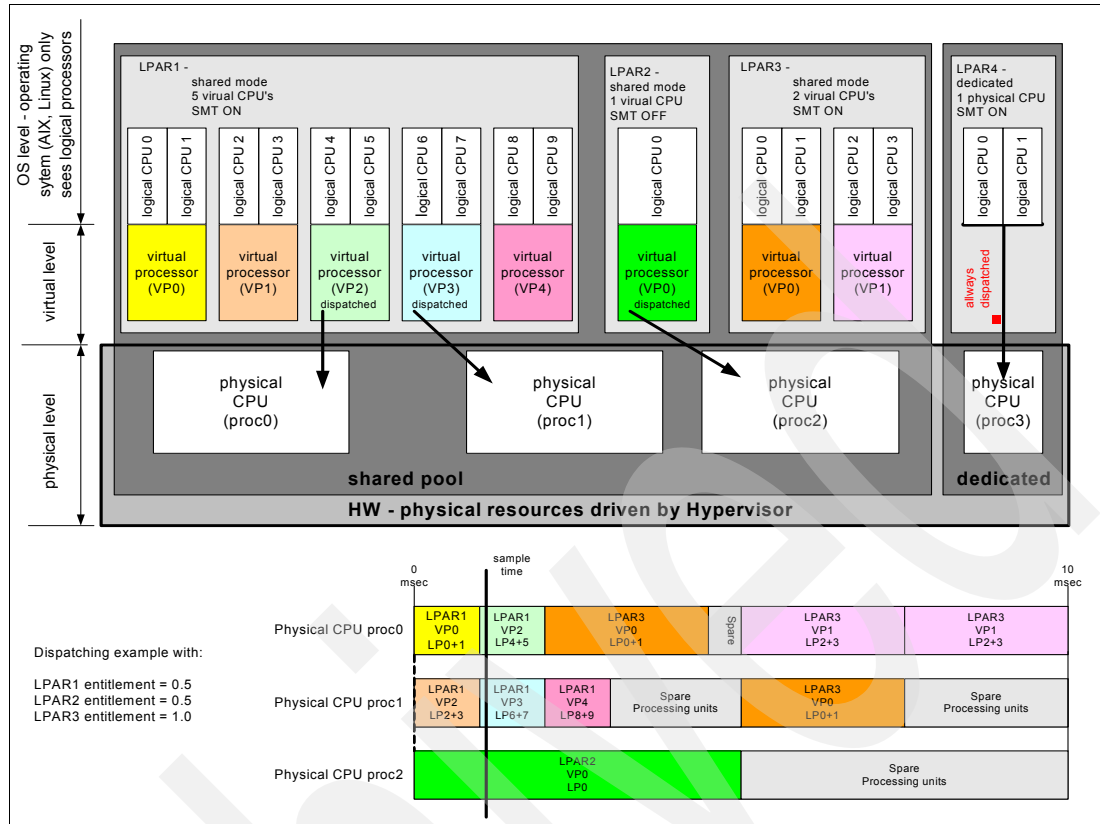


Figure 2-10 Logical, virtual, and physical processor mapping

In Figure 2-10, a system with four physical processors and four partitions is presented; one partition (LPAR4) is in dedicated mode and three partitions (LPAR1, LPAR2, and LPAR3) are running in shared mode. Dedicated mode LPAR4 is using one physical processor and thus three processors are available for shared processor pool. LPAR1 defines five virtual processors and the simultaneous multithreading feature is on (thus sees 10 logical processors), LPAR2 defines one virtual processor and simultaneous multithreading is off (one logical processor). LPAR3 defines two virtual processors and simultaneous multithreading is on. Currently (sample time), virtual processors 2 and 3 of LPAR1 and virtual processor 0 of LPAR2 are dispatched on physical processors in the shared pool. Other virtual processors are idle waiting for dispatch by the Hypervisor. When more virtual processors are defined within a partition, any virtual processor share equal parts of partition processing entitlement.

2.10.3 Virtual I/O Server

The Virtual I/O Server is a special purpose partition that provides virtual I/O resources to other partitions. The Virtual I/O Server owns the physical resources (SCSI, Fibre Channel, and network adapters, and optical devices) and allows client partitions to share access to them, thus minimizing the number of physical adapters in the system. The Virtual I/O Server eliminates the requirement that every partition own a dedicated network adapter, disk adapter, and disk drive.

Figure 2-11 on page 47 shows an organization view of a micro-partitioned system including the Virtual I/O Server. The figure also includes virtual SCSI and Ethernet connections and mixed operating system partitions.




Figure 2-11 Micro-Partitioning technology and VIOS

Because the Virtual I/O server is an operating system-based appliance server, redundancy for physical devices attached to the Virtual I/O Server can be provided by using capabilities such as Multipath I/O and IEEE 802.3ad Link Aggregation.

Installation of the Virtual I/O Server partition is performed from a special system backup DVD that is provided to clients that order the Advanced POWER Virtualization feature. This dedicated software is only for the Virtual I/O Server (and IVM in case it is used) and is only supported in special Virtual I/O Server partitions.

The Virtual I/O Server can be installed by:

- ▶ Media (assigning the DVD-ROM drive to the partition and booting from the media)
- ▶ The HMC (inserting the media in the DVD-ROM drive on the HMC and using the `installios` command)
- ▶ Using the Network Install Manager (NIM)

Note: To increase the performance of I/O-intensive applications, use dedicated physical adapters that use dedicated partitions.

We recommend that you install the Virtual I/O Server in a partition with dedicated resources or at least 0.5 processor entitlement to help ensure consistent performance.

The Virtual I/O Server supports RAID configurations and SAN attached devices (possibly with multipath driver). Logical volumes created on RAID or JBOD configurations are bootable, and the number of logical volumes is limited to the amount of storage available and architectural limits of the Logical Volume Manager.

Two major functions are provided with the Virtual I/O Server: a shared Ethernet adapter and Virtual SCSI.

Shared Ethernet adapter

A shared Ethernet adapter (SEA) is a Virtual I/O Server service that acts as a layer 2 network bridge between a physical Ethernet adapter or aggregation of physical adapters

(EtherChannel) and one or more virtual Ethernet adapters defined by Hypervisor on the Virtual I/O Server. A SEA enables LPARs on the virtual Ethernet to share access to the physical Ethernet and communicate with standalone servers and LPARs on other systems. The shared Ethernet network provides this access by connecting the internal Hypervisor VLANs with the VLANs on the external switches. Because the shared Ethernet network processes packets at layer 2, the original MAC address and VLAN tags of the packet is visible to other systems on the physical network. IEEE 802.1 VLAN tagging is supported.

The virtual Ethernet adapters that are used to configure a shared Ethernet adapter are required to have the trunk setting enabled. The trunk setting causes these virtual Ethernet adapters to operate in a special mode so that they can deliver and accept external packets from the POWER5 internal switch to the external physical switches. The trunk setting should only be used for the virtual Ethernet adapters that are part of a shared Ethernet setup in the Virtual I/O Server.

A single SEA setup can have up to 16 Virtual Ethernet trunk adapters and each virtual Ethernet trunk adapter can support up to 20 VLAN networks. Therefore, it is possible for a single physical Ethernet to be shared between 320 internal VLAN. The number of shared Ethernet adapters that can be set up in a Virtual I/O server partition is limited only by the resource availability as there are no configuration limits.

For a more detailed discussion about virtual networking, see:

http://www.ibm.com/servers/aix/whitepapers/aix_vn.pdf

Virtual SCSI

Access to real storage devices is implemented through the virtual SCSI services, a part of the Virtual I/O Server partition. This is accomplished using a pair of virtual adapters: a virtual SCSI server adapter and a virtual SCSI client adapter. The virtual SCSI server and client adapters are configured using an HMC or through Integrated Virtualization Manager on smaller systems. The virtual SCSI server (target) adapter is responsible for executing any SCSI commands it receives. It is owned by the Virtual I/O Server partition. The virtual SCSI client adapter allows a client partition to access physical SCSI and SAN attached devices and LUNs that are assigned to the client partition.

Physical disks owned by the Virtual I/O Server partition can either be exported and assigned to a client partition as whole device, or can be configured into a volume group and partitioned into several logical volumes. These logical volumes can then be assigned to individual partitions. From client partition point of view these two options are equivalent.

The Virtual I/O Server provides mapping between *backing devices* (physical devices or logical volumes assigned to client partitions in VIOS nomenclature) and client partitions by a command line interface. The appropriate command is the `mkvdev` command. For syntax and semantics see the Virtual I/O server documentation.

All current storage device types, such as SAN, SCSI, and RAID are supported, SSA and iSCSI are not supported at the time of writing.

For more information about the specific storage devices supported, see:

<http://techsupport.services.ibm.com/server/vios/home.html>

Note: Mirrored Logical Volumes (LVs) on Virtual I/O Server level are not recommended as backing devices. If mirroring is required, two independent devices (possibly from two separate VIO servers) should be assigned to the client partition and client partition should define mirror on top of them.

Virtual I/O Server Version 1.3

Virtual I/O Server Version 1.3 brings a host of new enhancements, including improved monitoring, such as additional **topas** and **viostat** performance metrics, and the bundling of the Performance ToolKit (PTX®) agent. Virtual SCSI and Virtual Ethernet performance increases, and command-line enhancements and enablement of additional storage solutions are also included.

Virtual I/O Server V1.3 introduced several enhancements for Virtual SCSI and shared Fibre Channel adapters support:

- ▶ Independent Software Vendor / Independent Hardware Vendor Virtual I/O enablement
- ▶ iSCSI TOE adapter
- ▶ iSCSI direct attached n3700 storage subsystem
- ▶ HP storage
- ▶ Virtual SCSI functional enhancements
 - Support for SCSI Reserve/Release for limited configurations
 - Changeable queue depth
 - Enables updating virtual device capacity non disruptively so that the virtual disk can “grow” without requiring a reconfiguration
 - Configurable fast fail time (number of retries on failure)
 - Error log enhancements

Virtual I/O Server V1.3 also introduced several enhancements for Virtual Ethernet and shared Ethernet adapter support.

2.10.4 Partition Load Manager

Partition Load Manager (PLM) provides automated processor and memory distribution between a dynamic LPAR and a Micro-Partitioning technology-capable logical partition running AIX 5L. The PLM application is based on a client/server model to share system information, such as processor or memory events, across the concurrent present logical partitions.

The following events are registered on all managed partition nodes:

- ▶ Memory-pages-steal high thresholds and low thresholds
- ▶ Memory-usage high thresholds and low thresholds
- ▶ Processor-load-average high threshold and low threshold

Note: PLM is supported on AIX 5L Version 5.2 and AIX 5L Version 5.3; it is not supported on Linux.

2.10.5 Integrated Virtualization Manager

In order to ease virtualization technology adoption in any IBM System p5 environment, IBM has developed Integrated Virtualization Manager (IVM), a simplified hardware management solution that inherits some HMC features, thus avoiding the necessity of a dedicated control workstation. This solution enables the administrator to reduce system setup time. IVM is targeted at small and medium systems.

IVM supports up to the maximum 16-core configuration. The IVM provides a management model for a single system. Although it does not provide the full flexibility of an HMC, it enables

the exploitation of the IBM Virtualization Engine™ technology. IVM is an enhancement of Virtual I/O Server offered as part of Virtual I/O Server Version 1.2 and follow on versions (the product that enables I/O virtualization in POWER5 and POWER5+ systems). It provides the same Virtual I/O Server features plus a Web-based graphical interface that enables the administrator to remotely manage the System p5 server with an Internet browser.

IVM can be used to complete the following tasks:

- ▶ Create and manage logical partitions.
- ▶ Configure the virtual Ethernet networks.
- ▶ Manage storage in the Virtual I/O Server.
- ▶ Create and manage user accounts.
- ▶ Create and manage serviceable events through Service Focal Point.
- ▶ Download and install updates to device microcode and to Virtual I/O Server software.
- ▶ Back up and restore logical partition configuration information.
- ▶ View application logs and the device inventory.

The requirements for an IVM managed server are as follows:

- ▶ A server managed by IVM cannot be simultaneously managed by an HMC.
- ▶ IVM (with Virtual I/O Server) must be installed as the first operating system.
- ▶ An IVM partition requires a minimum of one virtual processor and 512 MB of RAM.

Virtual I/O Server Version 1.3 introduced enhancements to IVM, which include support for Dynamic Logical Partitioning (DLPAR) for memory and processors in managed partitions. Additionally, a number of usability enhancements include support through the browser-based interface for IP configuration of the Virtual I/O Server:

- ▶ Dynamic LPAR support for memory and processors in managed partitions
- ▶ GUI support for System Plan management, including the Logical Partition (LPAR) Deployment Wizard
- ▶ Web UI support for:
 - IP configuration support
 - Task Manager for long-running tasks
 - Various usability enhancements, including the ability to create a new partition based on an existing one

The major considerations of IVM in comparison to an HMC-managed system are as follows:

- ▶ All physical adapters are owned by IVM, and LPARs use virtual devices only.
- ▶ There is only one profile per partition.
- ▶ A maximum of four virtual Ethernet networks are available inside the system.
- ▶ Each LPAR can have a maximum of one Virtual SCSI adapter assigned.
- ▶ IVM supports a single Virtual I/O Server to support all your mission critical production needs.
- ▶ Service Agent (see 3.2.3, “Service Agent” on page 75) for reporting hardware errors to IBM is not available on IVM.
- ▶ IVM cannot be used by HACMP software to activate CoD resources on machines that support CoD.

Despite those considerations, IVM provides advanced virtualization functionality without the need for an extra-cost workstation. For more information about IVM functionality and best practices, see *Virtual I/O Server Integrated Virtualization Manager*, REDP-4061, found at: <http://www.ibm.com/systems/p/hardware/meetp5/ivm.pdf>

Figure 2-12 shows how a system with IVM is organized. There is a Virtual I/O server and IVM installed in one partition owning all physical server resources and four client partitions. IVM communicates to Hypervisor to *create, manage, and provide virtual I/O* for client partitions. The dispatch of partitions on physical processors is done by the POWER Hypervisor, as in HMC managed servers. The rules for mapping the physical processors, virtual processors and logical processors apply, as discussed in 2.10.2, “Logical, virtual, and physical processor mapping” on page 44, for shared partitions managed by the HMC.

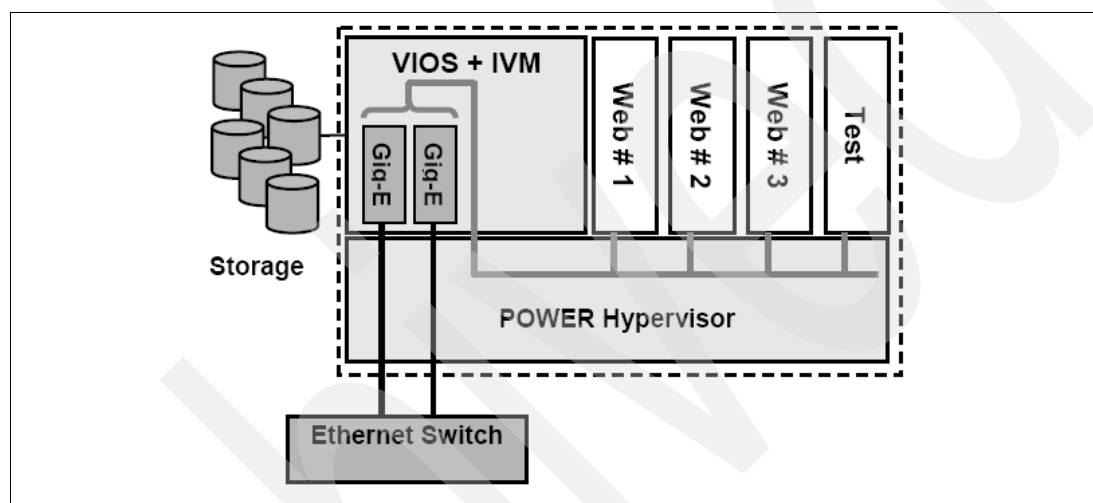


Figure 2-12 IVM principles

Note: IVM and HMC are two separate management systems and cannot be used at the same time. IVM targets ease of use, while HMC targets flexibility and scalability. The internal design is so different that you should never connect an HMC to a working IVM system. If you want to migrate an environment from IVM to HMC, you have to rebuild the configuration setup manually.

Operating system support for advanced virtualization

Table 2-10 lists AIX 5L and Linux support for advanced virtualization.

Table 2-10 Operating system supported functions

Advanced virtualization feature	AIX 5L Version 5.2	AIX 5L Version 5.3	Linux SLES 9	Linux RHEL AS 3	Linux RHEL AS 4
Micro-partitions (1/10th of processor)	N	Y	Y	Y	Y
Virtual Storage	N	Y	Y	Y	Y
Virtual Ethernet	N	Y	Y	Y	Y
Partition Load Manager	Y	Y	N	N	N

2.11 Hardware Management Console

The HMC is a dedicated workstation that provides a graphical user interface for configuring, operating, and performing basic system tasks for the System p5 servers functioning in either non-partitioned, LPAR, or clustered environments. In addition the Hardware Management Console is used to configure and manage partitions. One HMC is capable of controlling multiple POWER5 and POWER5+ processor-based systems.

At the time of writing, one HMC supports up to 48 POWER5 and POWER5+ processor-based systems and up to 254 LPARs using the HMC machine code Version 5.1. For updates of the machine code and HMC functions and hardware prerequisites, refer to the following Web site:

<http://techsupport.services.ibm.com/server/hmc>

POWER5 and POWER5+ processor-based system HMCs require Ethernet connectivity between HMC and server's service processor, moreover if dynamic LPAR operations are required, all AIX 5L and Linux partitions must be enabled to communicate over network to HMC. Ensure that sufficient Ethernet adapters are available to enable public and private networks, if you need both:

- ▶ The HMC 7310 Model C05 is a desktside model with only one integrated 10/100/1000 Mbps Ethernet port, but two additional PCI slots.
- ▶ The 7310 Model CR4 is a 1U, 19-inch rack-mountable drawer that has two native 10/100/1000 Mbps Ethernet ports and two additional PCI slots.

For any partition in a server, it is possible to use the shared Ethernet adapter in Virtual I/O Server for a unique connection from HMC to partitions. Therefore client partitions do not require own physical adapter in order to be able to communicate to HMC.

It is a good practice to connect the HMC to the first HMC port on the system, labeled as HMC Port 1, although other network configurations are possible. A second HMC can be attached to HMC Port 2 of the server for redundancy (or vice versa). Figure 2-13 shows a simple network configuration to enable the connection from HMC to server, and to enable dynamic LPAR operations. For more details about the HMC and the possible network connections, refer to:

<http://www.redbooks.ibm.com/abstracts/redp3999.html>

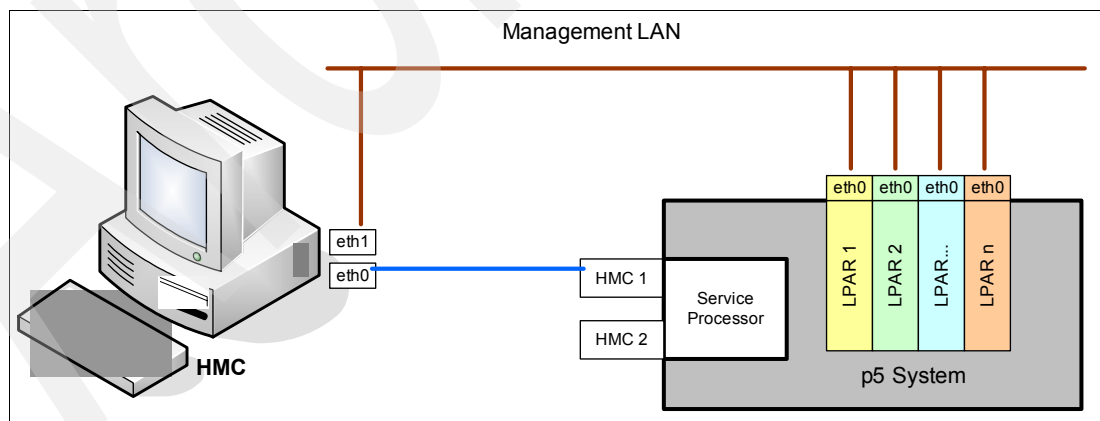


Figure 2-13 HMC to service processor and LPARs network connection

The default mechanism for allocation of the IP addresses for the service processor HMC ports is dynamic. The HMC can be configured as a DHCP server, providing the IP address at the time the managed server is powered on. If service processor of the managed server does not receive DHCP reply before time-out, predefined IP addresses will setup on both ports. Static IP address allocation is also an option. You can configure the IP address of the service

processor ports with a static IP address by using the Advanced System Management Interface (ASMI) menus. See 2.13.4, “Service processor” on page 63 for predefined IP addresses and additional information.

Note: If you need to access ASMI (for example to setup IP address of a new POWER5+ processor-based server when HMC is not available or not providing DHCP services), you can connect any client to one of the service processor HMC ports with any kind of Ethernet cable, and use a Web browser to access the predefined IP address, such as the following example:

`https://192.168.2.147`

Functions performed by the HMC include:

- ▶ Creating and maintaining a multiple partition environment
- ▶ Displaying a virtual operating system session terminal for each partition
- ▶ Displaying a virtual operator panel of contents for each partition
- ▶ Detecting, reporting, and storing changes in hardware conditions
- ▶ Powering managed systems on and off
- ▶ Acting as a service focal point

The HMC provides both graphical and command line interface for all management tasks. Remote connection to the HMC using Web-based System Manager or SSH are possible. For accessing the graphical interface, you can use the Web-based System Manager Remote Client running on the AIX 5L, Linux, or Windows® operating system. The Web-based System Manager client installation image can be downloaded from the HMC itself from the following URL:

`http://<hmc_address_or_name>/remote_client.html`

Both un-encrypted and encrypted Web-based System Manager connections are supported. The command line interface is also available by using the SSH secure shell connection to the HMC. It can be used by an external management system or a partition to perform HMC operations remotely.

2.11.1 High availability using the HMC

The HMC is an important hardware component. HACMP Version 5.3 High Availability cluster software can be used to automatically activate resources (where available) thus becoming an integral part of the cluster. For some environments, it is recommended to work with redundant HMCs. POWER5 and POWER5+ processor-based systems have two service processor interfaces (HMC port 1 and HMC port 2) available for connection to the HMC. It is recommended to use both of them for redundant network configuration. Depending on your environment, you have multiple options to configure the network. Figure 2-14 shows one possible highly available configuration.




Figure 2-14 Highly available HMC and network architecture

Note that only hardware management network (LAN1 and LAN2) is highly available on the above picture in order to keep simplicity. But also management network (LAN3) can be made highly available by using similar concept and adding more Ethernet adapters to LPARs and HMCs.

2.11.2 IBM System Planning Tool

The IBM System Planning Tool (SPT) is the next generation of the IBM LPAR Validation Tool (LVT). It contains all of the functions from the LVT and is integrated with the IBM Systems Workload Estimator (WLE). System plans generated by the SPT can be deployed on the system by the Hardware Management Console (HMC). The SPT is available to assist the user in system planning, design, validation, and to provide a system validation report that reflects the user's system requirements while not exceeding system recommendations. The SPT is a PC-based browser application designed to be run in a stand-alone environment.

The IBM System Planning Tool can be downloaded at no additional charge from:

<http://www.ibm.com/servers/eserver/support/tools/systemplanningtool/>

The System Planning Tool (SPT) helps you design a system to fit your needs. You can use the SPT to design a logically partitioned system or you can use the SPT to design an unpartitioned system. You can create an entirely new system configuration, or you can create a system configuration based upon any of the following:

- ▶ Performance data from an existing system that the new system is to replace
- ▶ Performance estimates that anticipate future workloads that you must support
- ▶ Sample systems that you can customize to fit your needs

Integration between the SPT and both the Workload Estimator (WLE) and IBM Performance Management (PM) allows you to create a system that is based upon performance and capacity data from an existing system or that is based on new workloads that you specify.

You can use the SPT before you order a system to determine what you must order to support your workload. You can also use the SPT to determine how you can partition a system that you already have.

Important: We recommend using the IBM System Planning Tool to estimate Hypervisor requirements and determine the memory resources that are required for all partitioned and non-partitioned servers.

Figure 2-15 shows the estimated Hypervisor memory requirements based on sample partition requirements.

Name	ID	OS type	Virtual memory (MB)			Virtual adapter count				
			Min	Desired	Max	Client serial	Ethernet	Client SCSI	Server SCSI	Reserved
* LPAR1	1	AIX_53	128	128	128	0	1	0	0	7
LPAR2	2	Linux_Virtual_Client	128	128	128	0	1	1	0	6
LPAR3	3	AIX_Virtual_Client	128	128	128	0	1	1	0	6
LPAR4	4	Virtual I/O Server	128	128	128	6	1	0	3	0
LPAR5	5	AIX_53	128	128	128	0	1	0	0	7
LPAR6	6	AIX_52	128	128	128	0	0	0	0	0
LPAR7	7	AIX_Virtual_Client	128	128	128	0	1	1	0	6

Figure 2-15 IBM System Planning Tool window showing Hypervisor requirements

2.12 Operating system support

The p5-560Q is capable of running the AIX 5L and Linux operating systems. The AIX 5L operating system has been specifically developed and enhanced to exploit and support the extensive RAS features on IBM System p systems.

2.12.1 AIX 5L

If installing AIX 5L on the p5-560Q with the 1.5 GHz processor option, the following minimum requirements must be met:

- ▶ AIX 5L for POWER V5.2 with 5200-08 Technology Level (APAR IY77270), CD# LCD4-1133-08 or later, DVD LCD4-7549-01 media is also available.
- ▶ AIX 5L for POWER V5.3 with 5300-04 Technology Level (APAR IY77273), CD#LCD4-7463-05 or later, DVD LCD4-7544-01 media is also available.

For the 1.8 GHz processor, the following levels are required:

- ▶ AIX 5L for POWER V5.2 with 5200-09 Technology Level (APAR IY82425), or later.
- ▶ AIX 5L for POWER V5.3 with 5300-05 Technology Level (APAR IY82426), or later.

Note: The Advanced POWER Virtualization feature (FC 7304) is not supported on AIX 5L V5.2; it requires AIX 5L V5.3.

IBM releases maintenance packages for the AIX 5L operating system periodically. These packages are available on CD-ROM, or you can download them from:

<http://www.ibm.com/servers/eserver/support/pseries/aixfixes.html>

The Web page provides information about how to obtain the CD-ROM.

You can also get individual operating system fixes and information about obtaining AIX 5L service at this site. In AIX 5L V5.3, the **suma** command is also available, which helps the administrator to automate the task of checking and downloading operating system downloads. For more information about the **suma** command functionality, refer to:

<http://techsupport.services.ibm.com/server/suma/home.html>

On 18 November 2005, Electronic Software Delivery (ESD) for AIX 5L V5.2 and V5.3 for POWER5 systems was made available. This delivery method is a way for clients to receive software and associated publications online, instead of waiting for a physical shipment to arrive. Clients requesting ESD should order FC 3450.

ESD has the following requirements:

- ▶ POWER5 system
- ▶ Internet connectivity from a POWER5 system or PC, reasonable connection speed for downloading large products such as AIX 5L
- ▶ Registration on the ESD Web site

For additional information, contact your IBM sales representative.

Software support for new features in POWER5+ processor

For a complete list of new features introduced in POWER5+ processor, see 2.1, “The POWER5+ processor” on page 24. Support for two new virtual memory page sizes was introduced - 64 KB and 16 GB as well as support for 1 TB segment size. While 16 GB pages are intended to only be used in very high performance environments, 64 KB pages are general-purpose. AIX 5L Version 5.3 with the 5300-04 Technology Level 64-bit kernel is required for 64 KB and 16 GB page size support.

As with all previous versions of AIX 5L, 4 KB is the default page size. A process will continue to use 4 KB pages unless a user specifically requests another page size be used. AIX 5L has rich support of 64 KB pages. They are easy to use, and it is expected that many applications will see performance benefits when using 64 KB pages rather than 4 KB pages. No system configuration changes are necessary to enable a system to use 64 KB pages, they are fully pageable, and the size of the pool of 64 KB page frames on a system is dynamic and fully managed by AIX 5L.

The main benefit of a larger page size is improved performance for applications that allocate and repeatedly access large amounts of memory. The performance improvement from larger page sizes is due to the overhead of translating a page address as it is used in an application, to a page address that is understood by the computer's memory subsystem. To improve performance, the information needed to translate a given page is usually cached in the processor. In POWER5+, this cache takes the form of a translation lookaside buffer (TLB). Since there are a limited number of TLB entries, using a large page size increases the amount of address space that can be accessed without incurring translation delays. Also the size of TLB in POWER5+ has been doubled compared to POWER5.

Huge pages (16 GB) are intended to be used only in very high performance environments, and AIX 5L will not automatically configure a system to use these page sizes. A system administrator must configure AIX 5L to use these page sizes and specify their number via HMC before partition start.

A user can specify page sizes to use for three regions process's address space with an environment variable or with settings in an application's XCOFF binary using the `ldedit` or `ld` commands. These three regions are: data, stack and program text. An application programmer can also select the page size to use for System V shared memory via a new `SHM_PAGESIZE` command to the `shmctl()` system call.

An example of using system variables to start a program with 64 KB page size support:

```
LDR_CNTRL=DATAPSIZE=64K@TEXTPSIZE=64K@STACKPSIZE=64K <program>
```

Systems commands (`ps`, `vmstat`, `svmon`, `pagesize`) have been enhanced to report various page size usage.

2.12.2 Linux

For the p5-560Q, Linux distributions are available through Novel SUSE and Red Hat at the time this publication was written. The p5-560Q requires the following version of Linux distributions:

- ▶ SUSE Linux Enterprise Server 9 for IBM POWER Service Pack 3, or later
- ▶ Red Hat Enterprise Linux AS 4 for IBM POWER Service Update 2, or later

Note: Not all p5-560Q features available on AIX 5L are available on Linux.

For information about the features and external devices supported by Linux, refer to:

<http://www.ibm.com/servers/eserver/pseries/linux/>

For information about SUSE Linux Enterprise Server 9, refer to:

<http://www.novell.com/products/linuxenterpriseserver/>

For information about Red Hat Enterprise Linux AS, refer to:

<http://www.redhat.com/software/rhel/details/>

Many of the features described in this document are operating system dependant and might not be available on Linux. For more information, see:

http://www.ibm.com/servers/eserver/linux/power/whitepapers/linux_overview.html

Note: IBM only supports the Linux systems of clients with a SupportLine contract covering Linux. Otherwise, contact the Linux distributor for support.

Specially-priced Linux subscriptions

Linux subscriptions are now available when ordered through IBM and combined with an IBM System p5 Express Product Offering. Clients can purchase a one-year specially priced subscription or a greater discount for a three-year subscription.

These new Linux options, available on System p5 Express servers, bring improved pricing and price performance to our clients interested in Linux as their primary operating system. Clients interested in AIX 5L can also obtain an Express Product Offering that fits their needs.

Clients are still encouraged to purchase support for their Linux subscription either through IBM Global Services or through the distributor to receive updates and technical assistance as needed. Support is not included in the price of the subscription.

The new lower-priced Linux subscriptions, when combined with the lower package prices of the System p5 Express Product Offering, make these products an exceptional value for our smaller to mid-market clients, as well as larger enterprises.

Refer to the following Web site for Red Hat information:

<http://www.redhat.com/software/>

For additional information about Linux on POWER, visit:

<http://www.ibm.com/servers/eserver/linux/power/>

2.13 Service information

The p5-560Q is not customer setup server (CSU). Therefore, the IBM service representative completes the system installation.

2.13.1 Touch point colors

Blue (IBM blue) or terra-cotta (orange) on a component indicates a touch point (for electronic parts) where you can grip the hardware to remove it from or to install it into the system, to open or to close a latch, and so on. IBM defines the touch point colors as follows:

Blue

This requires a shutdown of the system before the task can be performed, for example, installing additional processors contained in the second processor book.

Terra-cotta

The system can remain powered on while this task is being performed. Keep in mind that some tasks might require that other steps have to be performed first. One example is deconfiguring a physical volume in the operating system before removing a disk from a 4-pack disk enclosure of the server.

Blue and terra-cotta

Terra-cotta takes precedence over this color combination, and the rules for a terra-cotta-only touch point apply.

Important: It is important to adhere to the touch point colors on the system. Not doing so can compromise your safety and damage the system.

2.13.2 Operator control panel

The service processor provides an interface to the control panel that is used to display server status and diagnostic information. See Figure 2-16 for operator control panel physical details and buttons.




Figure 2-16 Operator control panel physical details and button

Note: For servers managed by the HMC, use it to perform control panel functions.

Primary control panel functions

The primary control panel functions are defined as functions 01 to 20, including options to view and manipulate IPL modes, server operating modes, IPL speed, and IPL type.

The following list describes the primary functions:

- ▶ Function 01: Display selected IPL type, system operating mode, and IPL speed
- ▶ Function 02: Select IPL type, IPL speed override, and system operating mode
- ▶ Function 03: Start IPL
- ▶ Function 04: Lamp Test
- ▶ Function 05: Reserved
- ▶ Function 06: Reserved
- ▶ Function 07: SPCN functions
- ▶ Function 08: Fast Power Off
- ▶ Functions 09 to 10: Reserved
- ▶ Functions 11 to 19: System Reference Code
- ▶ Function 20: System type, model, feature code, and IPL type

All the functions mentioned are accessible using the Advanced System Management Interface (ASMI), HMC, or the control panel.

Extended control panel functions

The extended control panel functions consist of two major groups:

- ▶ Functions 21 through 49, which are available when you select Manual mode from Function 02.
- ▶ Support service representative Functions 50 through 99, which are available when you select Manual mode from Function 02, then select and enter the client service switch 1 (Function 25), followed by service switch 2 (Function 26).

Function 30 – CEC SP IP address and location

Function 30 is one of the Extended control panel functions and is only available when Manual mode is selected. This function can be used to display the central electronic complex (CEC) Service Processor IP address and location segment. The Table 2-11 shows an example of how to use the Function 03:

Table 2-11 CEC SP IP address and location

Information on operator panel	Action or description
3 0	Use the increment or decrement buttons to scroll to Function 30
3 0 * *	Press Enter to enter sub-function mode
3 0 0 0	Use the increment or decrement buttons to select an IP address: 0 0 = Service Processor ETH0 or HMC1 port 0 1 = Service Processor ETH1 or HMC2 port
S P A : E T H 0 : _ _ _ T 5 1 9 2 . 1 6 8 . 2 . 1 4 7	Press Enter to display the selected IP address
3 0 * *	Use the increment or decrement buttons to select sub-function exit
3 0	Press Enter to exit sub-function mode

2.13.3 System firmware

Server firmware is the part of the Licensed Internal Code that enables hardware, such as the service processor. Depending on your service environment, you can download, install, and manage your server firmware fixes using different interfaces and methods, including the HMC, or by using functions specific to your operating system. See 3.2.4, “IBM System p5 firmware maintenance” on page 77 for a detailed description of System p5 firmware.

Important: Prior to installing a 1.8 GHz processor feature, the firmware level must be at least SF240_261 or later.

Note: Normally, installing the server firmware fixes through the operating system is a nonconcurrent process.

Temporary and permanent firmware sides

The service processor maintains two copies of the server firmware:

- ▶ One copy is considered the permanent or backup copy and is stored on the permanent side, sometimes referred to as the *p* side.
- ▶ The other copy is considered the installed or temporary copy and is stored on the temporary side, sometimes referred to as the *t* side. We recommend that you start and run the server from the temporary side.

The copy actually booted from is called the activated level, sometimes referred to as *b*.

Note: The default value, from which the system boots, is temporary.

The following examples are the output of the `lsmcodes` command for AIX 5L and Linux, showing the firmware levels as they are displayed in the outputs.

▶ AIX 5L:

The current permanent system firmware image is SF220_005.
The current temporary system firmware image is SF220_006.
The system is currently booted from the temporary image.

► Linux:

system:SF220_006 (t) SF220_005 (p) SF220_006 (b)

When you install a server firmware fix, it is installed on the temporary side.

Note: The following points are of special interest:

- The server firmware fix is installed on the temporary side only after the existing contents of the temporary side are permanently installed on the permanent side (the service processor performs this process automatically when you install a server firmware fix).
- If you want to preserve the contents of the permanent side, you need to remove the current level of firmware (copy the contents of the permanent side to the temporary side) before you install the fix.
- However, if you get your fixes using the Advanced features on the HMC interface and you indicate that you do not want the service processor to automatically accept the firmware level, the contents of the temporary side are not automatically installed on the permanent side. In this situation, you do not need to remove the current level of firmware to preserve the contents of the permanent side before you install the fix.

You might want to use the new level of firmware for a period of time to verify that it works correctly. When you are sure that the new level of firmware works correctly, you can permanently install the server firmware fix. When you permanently install a server firmware fix, you copy the temporary firmware level from the temporary side to the permanent side.

Conversely, if you decide that you do not want to keep the new level of server firmware, you can remove the current level of firmware. When you remove the current level of firmware, you copy the firmware level that is currently installed on the permanent side from the permanent side to the temporary side.

System firmware download site

For the system firmware download site for the p5-560Q, go to:

<http://techsupport.services.ibm.com/server/mdownload>

Receive server firmware fixes using an HMC

If you use an HMC to manage your server and you periodically configured several partitions on the server, you need to download and install fixes for your server and power subsystem firmware.

How you get the fix depends on whether the HMC or server is connected to the Internet:

- The HMC or server is connected to the Internet.

There are several repository locations from which you can download the fixes using the HMC. For example, you can download the fixes from your service provider's Web site or support system, from optical media that you order from your service provider, or from an FTP server on which you previously placed the fixes.
- Neither the HMC nor your server is connected to the Internet (server firmware only).

You need to download your new server firmware level to a CD-ROM media or FTP server.

For both of these options, you can use the interface on the HMC to install the firmware fix (from one of the repository locations or from the optical media). The Change Internal Code wizard on the HMC provides a step-by-step process for you to perform the procedure to install the fix. Perform these steps:

1. Ensure that you have a connection to the service provider (if you have an Internet connection from the HMC or server).
2. Determine the available levels of server and power subsystem firmware.
3. Create optical media (if you do not have an Internet connection from the HMC or server).
4. Use the Change Internal Code wizard to update your server and power subsystem firmware.
5. Verify that the fix installed successfully.

For a detailed description of each task, select **Customer service, support, and troubleshooting** → **Fixes and upgrades** → **Getting fixes and upgrades** from the IBM Systems Hardware Information Center Web site at:

<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?lang=en>

Receive server firmware fixes without an HMC

Periodically, you need to install fixes for your server firmware. If you do not use an HMC to manage your server, you must get your fixes through your operating system. In this situation, you can get server firmware fixes through the operating system regardless of whether your operating system is AIX 5L or Linux.

To do this, complete the following tasks:

1. Determine the existing level of server firmware using the **lsmcode** command.
2. Determine the available levels of server firmware.
3. Get the server firmware.
4. Install the server firmware fix to the temporary side.
5. Verify that the server firmware fix installed successfully.
6. Install the server firmware fix permanently (optional).

Note: To view existing levels of server firmware using the **lsmcode** command, you need to have the following service tools installed on your server:

► AIX 5L

You must have AIX 5L diagnostics installed on your server to perform this task. AIX 5L diagnostics are installed when you install AIX 5L on your server. However, it is possible to deselect the diagnostics. Therefore, you need to ensure that the online AIX 5L diagnostics are installed before proceeding with this task.

► Linux

- Platform Enablement Library: `librtas-nnnnn.rpm`
- Service Aids: `ppc64-utils-nnnnn.rpm`
- Hardware Inventory: `lsvpd-nnnnn.rpm`

Where *nnnnn* represents a specific version of the RPM file.

If you do not have the service tools on your server, you can download them at the following Web site:

<http://techsupport.services.ibm.com/server/lopdiags>

See 3.2.4, “IBM System p5 firmware maintenance” on page 77 for additional information.

2.13.4 Service processor

The service processor is an embedded controller running the service processor internal operating system. The service processor operating system contains specific programs and device drivers for the service processor hardware. The host interface is a 32-bit PCI-X interface connected to the Enhanced I/O Controller.

Service processor is used to monitor and manage the system hardware resources and devices. The service processor offers the following connections:

- ▶ Two Ethernet 10/100 Mbps ports
 - Both Ethernet ports are only visible to the service processor and can be used to attach the p5-560Q to a HMC or to access the Advanced System Management Interface (ASMI) options from a client Web browser, using the HTTP server integrated into the service processor internal operating system.
 - Both Ethernet ports have a default IP address
 - Service processor Eth0 or HMC1 port is configured as 192.168.2.147
 - Service processor Eth1 or HMC2 port is configured as 192.168.3.147

2.13.5 Hardware management user interfaces

This section provides a brief overview of the different p5-560Q hardware management user interfaces available.

Advanced System Management Interface

The Advanced System Management Interface (ASMI) is the interface to the service processor that enables you to set flags that affect the operation of the server, such as auto power restart, and to view information about the server, such as the error log and vital product data.

This interface is accessible using a Web browser on a client system that is connected directly to the service processor (in this case a standard Ethernet cable or a crossed cable can be both used) or through an Ethernet network. Using the *network configuration menu*, the ASMI enables the possibility to change the service processor IP addresses or to apply some security policies and avoid the access from undesired IP addresses or range. The ASMI can also be accessed using a terminal attached to the system service processor ports on the server, if the server is not HMC managed. The service processor and the ASMI are standard on all IBM System p servers.

You might be able to use the service processor's default settings. In that case, accessing the ASMI is not necessary.

Accessing the ASMI using a Web browser

The Web interface to the Advanced System Management Interface is accessible through, at the time of writing, Microsoft® Internet Explorer® 6.0, Netscape 7.1, Mozilla Firefox, or Opera 7.23 running on a PC or mobile computer connected to the service processor. The Web interface is available during all phases of system operation including the initial program load and run time. However, some of the menu options in the Web interface are unavailable during IPL or run time to prevent usage or ownership conflicts if the system resources are in use during that phase.

Accessing the ASMI using an ASCII console

The Advanced System Management Interface on an ASCII console supports a subset of the functions provided by the Web interface and is available only when the system is in the

platform standby state. The ASMI on an ASCII console is not available during some phases of system operation, such as the initial program load and run time.

Accessing the ASMI using an HMC

To access the Advanced System Management Interface using the Hardware Management Console, complete the following steps:

1. Ensure that the HMC is set up and configured.
2. In the navigation area, expand the managed system with which you want to work.
3. Expand **Service Applications** and click **Service Focal Point**.
4. In the content area, click **Service Utilities**.
5. From the Service Utilities window, select the managed system with which you want to work with.
6. From the Selected menu on the Service Utilities window, select **Launch ASM menu**.

System Management Services

Use the System Management Services (SMS) menus to view information about your system or partition and to perform tasks such as changing the boot list, or setting the network parameters.

To start System Management Services, perform the following steps:

1. For a server that is connected to an HMC, use the HMC to restart the server or partition.
If the server is not connected to an HMC, stop the system, and then restart the server by pressing the power button on the control panel.
2. For a partitioned server, watch the virtual terminal window on the HMC.
For a full server partition, watch the firmware console.
3. Look for the POST³ indicators memory, keyboard, network, SCSI, speaker that appear across the bottom of the screen. Press the numeric 1 key after the word keyboard appears and before the word speaker appears.

The SMS menus is useful to define the operating system installation method, choosing the installation boot device or setting the boot device priority list for a full managed server or a logical partition. In case of a network boot, SMS menus provide to setup the network parameters and network adapter IP address.

HMC

The Hardware Management Console is a system that controls managed systems, including IBM System p5 hardware, and logical partitions. To provide flexibility and availability, there are different ways to implement HMCs.

Web-based System Manager Remote Client

The Web-based System Manager Remote Client is an application that is usually installed on a PC and can be downloaded directly from an installed HMC. When an HMC is installed and HMC Ethernet IP addresses have been assigned, it is possible to download the Web-based System Manager Remote Client from a web browser, using the following URL:

http://HMC_IP_address/remote_client.html

You can then use the PC to access other HMCs remotely. Web-based System Manager Remote Clients can be present in private and open networks. You can perform most

³ POST stands for power-on-self-test.

management tasks using the Web-based System Manager Remote Client. The remote HMC and the Web-based System Manager Remote Client allow you the flexibility to access your managed systems (including HMCs) from multiple locations using multiple HMCs.

For more detailed information about the use of the HMC, refer to the IBM Systems Hardware Information Center.

Open Firmware

A System p5 server has one instance of Open Firmware both when in the partitioned environment and when running as a full system partition. Open Firmware has access to all devices and data in the server. Open Firmware is started when the server goes through a power-on reset. Open Firmware, which runs in addition to the POWER Hypervisor in a partitioned environment, runs in two modes: global and partition. Each mode of Open Firmware shares the same firmware binary that is stored in the flash memory.

In a partitioned environment, Open Firmware runs on top of the global Open Firmware instance. The partition Open Firmware is started when a partition is activated. Each partition has its own instance of Open Firmware and has access to all the devices assigned to that partition. However, each instance of Open Firmware has no access to devices outside of the partition in which it runs. Partition firmware resides within the partition memory and is replaced when AIX 5L or Linux takes control. Partition firmware is needed only for the time that is necessary to load AIX 5L or Linux into the partition server memory.

The global Open Firmware environment includes the partition manager component. That component is an application in the global Open Firmware that establishes partitions and their corresponding resources (such as CPU, memory, and I/O slots), which are defined in partition profiles. The partition manager manages the operational partitioning transactions. It responds to commands from the service processor external command interface that originates in the application running on the HMC. The ASMI can be accessed during boot time or using the ASMI and selecting the boot to Open Firmware prompt.

For more information about Open Firmware, refer to *Partitioning Implementations for IBM eServer p5 Servers*, SG24-7039, which is available at:

<http://www.redbooks.ibm.com/abstracts/sg247039.html>

Archived



RAS and manageability

This chapter provides information about IBM System p5 design features that help lower the total cost of ownership (TCO). The state of art IBM RAS (Reliability, Availability, and Service ability) technology allows possibility to improve your TCO architecture by reducing unplanned down time. This chapter includes several features that are based on the benefits available when using AIX 5L. Support of these features using Linux can vary.

3.1 Reliability, availability, and serviceability

Excellent quality and reliability are inherent in all aspects of the IBM System p5 design and manufacturing. The fundamental objective of the design approach is to minimize outages. The RAS features help to ensure that the system operates when required, performs reliably, and efficiently handles any failures that might occur. This is achieved using capabilities provided by both the hardware and the operating system AIX 5L.

The p5-560Q as a POWER5+ processor-based server enhances the RAS capabilities implemented in POWER4 processor-based servers. RAS enhancements available are:

- ▶ Most firmware updates allow the system to remain operational.
- ▶ The ECC has been extended to inter-chip connections for the fabric and processor bus.
- ▶ Partial L2 cache deallocation is possible.
- ▶ The number of L3 cache line deletes improved from two to ten for better self-healing capability.

The following sections describe the concepts that form the basis of leadership RAS features of IBM System p5 product line in more details.

3.1.1 Fault avoidance

System p5 servers are built on a quality-based design intended to keep errors from happening. This design includes the following features:

- ▶ Reduced power consumption, cooler operating temperatures for increased reliability, enabled by the use of copper chip circuitry, silicon-on-insulator, and dynamic clock gating
- ▶ Mainframe-inspired components and technologies

3.1.2 First Failure Data Capture

If a problem should occur, the ability to correctly diagnose it is a fundamental requirement upon which improved availability is based. The p5-560Q incorporate advanced capability in start-up diagnostics and in run-time First Failure Data Capture (FDDC) that is based on strategic error checkers built into the chips.

Any errors that are detected by the pervasive error checkers are captured into Fault Isolation Registers (FIRs), which can be interrogated by the service processor. The service processor has the capability to access system components using special purpose ports or by access to the error registers. Figure 3-1 on page 69 shows a schematic of a Fault Register Implementation.

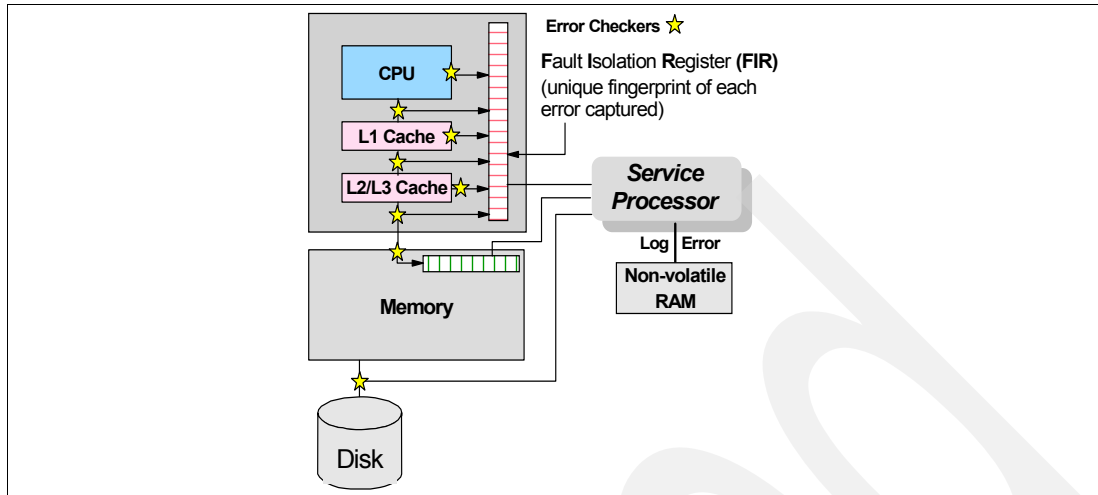


Figure 3-1 Schematic of Fault Isolation Register implementation

The FIRs are important because they enable an error to be uniquely identified, thus enabling the appropriate action to be taken. Appropriate actions might include such things as a bus retry, ECC correction, or system firmware recovery routines. Recovery routines can include dynamic deallocation of potentially failing components.

Errors are logged into the system non-volatile random access memory (NVRAM) and the service processor event history log, along with a notification of the event to AIX 5L for capture in the operating system error log. Diagnostic Error Log Analysis (*diagela*) routines analyze the error log entries and invoke a suitable action such as issuing a warning message. If the error can be recovered, or after suitable maintenance, the service processor resets the FIRs so that they can accurately record any future errors.

The ability to correctly diagnose any pending or firm errors is a key requirement before any dynamic or persistent component deallocation or any other reconfiguration can take place.

For further details, see 3.1.7, “Resource deallocation” on page 71.

3.1.3 Permanent monitoring

The service processor provides a way to monitor the system even when the main processor is inoperable.

Mutual surveillance

The service processor can monitor the operation of the firmware during the boot process, and it can monitor the operating system for loss of control. This allows the service processor to take appropriate action, including calling for service, when it detects that the firmware or the operating system has lost control. Mutual surveillance also allows the operating system to monitor for service processor activity and can request a service processor repair action if necessary.

Environmental monitoring

Environmental monitoring related to power, fans, and temperature is done by the System Power Control Network (SPCN). Environmental critical and non-critical conditions generate Early Power-Off Warning (EPOW) events. Critical events (for example, loss of primary power) trigger appropriate signals from hardware to impacted components so as to prevent any data

loss without the operating system or firmware involvement. Non-critical environmental events are logged and reported using Event Scan.

The operating system cannot program or access the temperature threshold using the service processor.

EPOW events can, for example, trigger the following actions.

- ▶ Temperature monitoring, which increases the fans speed rotation when ambient temperature is above a preset operating range.
- ▶ Temperature monitoring warns the system administrator of potential environmental-related problems. It also performs an orderly system shutdown when the operating temperature exceeds a critical level.
- ▶ Voltage monitoring provides warning and an orderly system shutdown when the voltage is out of the operational specification.

3.1.4 Self-healing

For a system to be self-healing, it must be able to recover from a failing component by first detecting and isolating the failed component, taking it offline, fixing or isolating it, and reintroducing the fixed or replacement component into service without any application disruption. Examples include:

- ▶ *Bit steering* to redundant memory in the event of a failed memory module to keep the server operational
- ▶ *Bit-scattering*, thus allowing for error correction and continued operation in the presence of a complete chip failure (*Chipkill™* recovery)
- ▶ There is ECC on the data received on the cache chip from the processor, which protects the interface for data from the processor to the cache
- ▶ There is ECC on the data read out of the eDRAM, which flags an array error.
- ▶ There is ECC on the processor receive interface, which protects the interface for data from the cache to the processor
- ▶ L3 cache line deletes extended from 2 to 10 for additional self-healing
- ▶ ECC extended to inter-chip connections on fabric and processor bus
- ▶ *Memory scrubbing* to help prevent soft-error memory faults

Memory reliability, fault tolerance, and integrity

The p5-560Q use Error Checking and Correcting (ECC) circuitry for system memory to correct single-bit and to detect double-bit memory failures. Detection of double-bit memory failures helps maintain data integrity. Furthermore, the memory chips are organized such that the failure of any specific memory module only affects a single bit within a four-bit ECC word (*bit-scattering*), thus allowing for error correction and continued operation in the presence of a complete chip failure (*Chipkill recovery*). The memory DIMMs also use *memory scrubbing* and thresholding to determine when spare memory modules within each bank of memory should be used to replace ones that have exceeded their threshold of error count (*dynamic bit-steering*). Memory scrubbing is the process of reading the contents of the memory during idle time and checking and correcting any single-bit errors that have accumulated by passing the data through the ECC logic. This function is a hardware function on the memory controller chip and does not influence normal system memory performance.

3.1.5 N+1 redundancy

The use of redundant parts allows the p5-560Q to remain operational with full resources:

- ▶ Redundant spare memory bits in L1, L2, L3, and main memory
- ▶ Redundant fans
- ▶ Redundant power supplies (optional)

Note: With this standard feature, every p5-560Q building block requires two power cords, which are not included in the base order. For maximum availability it is highly recommended to connect power cords from same p5-560Q building block to two separate PDUs in the rack. These PDUs being connected to two independent client power sources.

3.1.6 Fault masking

If corrections and retries succeed and do not exceed threshold limits, the system remains operational with full resources, and no intervention is required:

- ▶ CEC bus retry and recovery
- ▶ PCI-X bus recovery
- ▶ ECC Chipkill soft error

3.1.7 Resource deallocation

If recoverable errors exceed threshold limits, resources can be deallocated with the system remaining operational, allowing deferred maintenance at a convenient time.

Dynamic or persistent deallocation

Dynamic deallocation of potentially failing components is nondisruptive, allowing the system to continue to run. Persistent deallocation occurs when a failed component is detected, which is then deactivated at a subsequent reboot.

Dynamic deallocation functions include:

- ▶ Processor
- ▶ L3 cache line delete
- ▶ Partial L2 cache deallocation
- ▶ PCI-X bus and slots

For dynamic processor deallocation, the service processor performs a predictive failure analysis based on any recoverable processor errors that have been recorded. If these transient errors exceed a defined threshold, the event is logged and the processor is deallocated from the system while the operating system continues to run. This feature (named *CPU Guard*) enables maintenance to be deferred until a suitable time. Processor deallocation can only occur if there are sufficient functional processors (at least two).

To verify whether CPU Guard has been enabled, run the following command:

```
lsattr -El sys0 | grep cpuguard
```

If enabled, the output will be similar to the following:

```
cpuguard    enable    CPU Guard    True
```

If the output shows CPU Guard as disabled, enter the following command to enable it:

```
chdev -l sys0 -a cpuguard='enable'
```

Cache or cache-line deallocation is aimed at performing dynamic reconfiguration to bypass potentially failing components. This capability is provided for both L2 and L3 caches. Dynamic run-time deconfiguration is provided if a threshold of L1 or L2 recovered errors is exceeded.

In the case of an L3 cache run-time array single-bit solid error, the spare chip resources are used to perform a line delete on the failing line.

PCI-X hot-plug slot fault tracking helps prevent slot errors from causing a system machine check interrupt and subsequent reboot. This provides superior fault isolation, and the error affects only the single adapter. Run-time errors on the PCI bus caused by failing adapters will result in recovery action. If this is unsuccessful, the PCI device will be gracefully shut down. Parity errors on the PCI bus itself will result in bus retry, and if uncorrected, the bus and any I/O adapters or devices on that bus will be deconfigured.

The p5-560Q supports PCI Extended Error Handling (EEH) if it is supported by the PCI-X adapter. In the past, PCI bus parity errors caused a global machine check interrupt, which eventually required a system reboot in order to continue. In the p5-560Q system, hardware, system firmware, and AIX 5L interaction have been designed to allow transparent recovery of intermittent PCI bus parity errors and graceful transition to the I/O device available state in the case of a permanent parity error in the PCI bus.

EEH-enabled adapters respond to a special data packet generated from the affected PCI-X slot hardware by calling system firmware, which will examine the affected bus, allow the device driver to reset it, and continue without a system reboot.

Persistent deallocation functions include:

- ▶ Processor
- ▶ Memory
- ▶ Deconfigure or bypass failing I/O adapters
- ▶ L3 cache

Following a hardware error that has been flagged by the service processor, the subsequent reboot of the system will invoke extended diagnostics. If a processor or L3 cache has been marked for deconfiguration by persistent processor deallocation, the boot process will attempt to proceed to completion with the faulty device automatically deconfigured. Failing I/O adapters will be deconfigured or bypassed during the boot process.

Note: The auto-restart (reboot) option, when enabled, can reboot the system automatically following an unrecoverable software error, software hang, hardware failure, or environmentally induced failure (such as loss of power supply).

3.1.8 Serviceability

Increasing service productivity means the system is up and running for a longer time. p5-560Q improve service productivity by providing the functions described in the following sections:

Error indication and LED indicators

The p5-560Q is not designed for client setup and most of the hardware features can be handled by the IBM service representative. You might like to contact the IBM service representative to know which are the feature entitled to be a Client Replaceable Unit (CRU). The p5-560Q provides internal and external LED diagnostics that will identify parts that require service. Attenuation of the error is provided through a series of light attention signals,

starting on the exterior of the system (System Attention LED) located on the front of the system, and ending with an LED near the failing unit.

For more information about Client Replaceable Units, including videos, see:

<http://publib.boulder.ibm.com/eserver>

System Attention LED

The attention indicator is represented externally by an amber LED on the operator panel and the back of the system unit. It is used to indicate that the system is in one of the following states:

- ▶ Normal state, LED is off.
- ▶ Fault state, LED is on solid.
- ▶ Identify state, LED is blinking.

Additional LEDs on I/O components such as PCI-X slots and disk drives provide status information such as power, hot-swap, and need for service.

Concurrent Maintenance

Concurrent Maintenance provides replacement of the following parts while the system remains running:

- ▶ Disk drives
- ▶ Cooling fans
- ▶ Power subsystems
- ▶ PCI-X adapter cards
- ▶ Operator Panel (requires HMC guided support)

Remember to take in consideration the touch point colors, as described in 2.13.1, “Touch point colors” on page 58.

3.2 Manageability

The functions and tools provided for IBM System p5 servers to ease management are described in the next sections.

3.2.1 Service processor

The Service processor (SP) is always working regardless of main p5 Central Electronic Complex (CEC) state. CEC can be in the following states:

- ▶ Power standby mode (power off).
- ▶ Operating, ready to start partitions
- ▶ Operating with some partitions running and an AIX 5L or Linux system in control of the machine.

The SP is still working and checking the system for errors, ensuring the connection to the HMC (if present) for manageability purposes and accepting Advanced System Management Interface (ASMI) SSL network connections. The SP provides the possibility to view and manage the machine-wide settings using the ASMI and allows complete system and partition management from HMC. Also, the surveillance function of the SP is monitoring the operating system to check that it is still running and has not stalled.

Note: The IBM System p5 service processor enables the analysis of a system that will not boot. It can be performed either from ASMI, HMC or ASCI console (depending on presence of HMC). ASMI is provided in any case.

See Figure 3-2 for an example of the ASMI accessed from a Web browser.

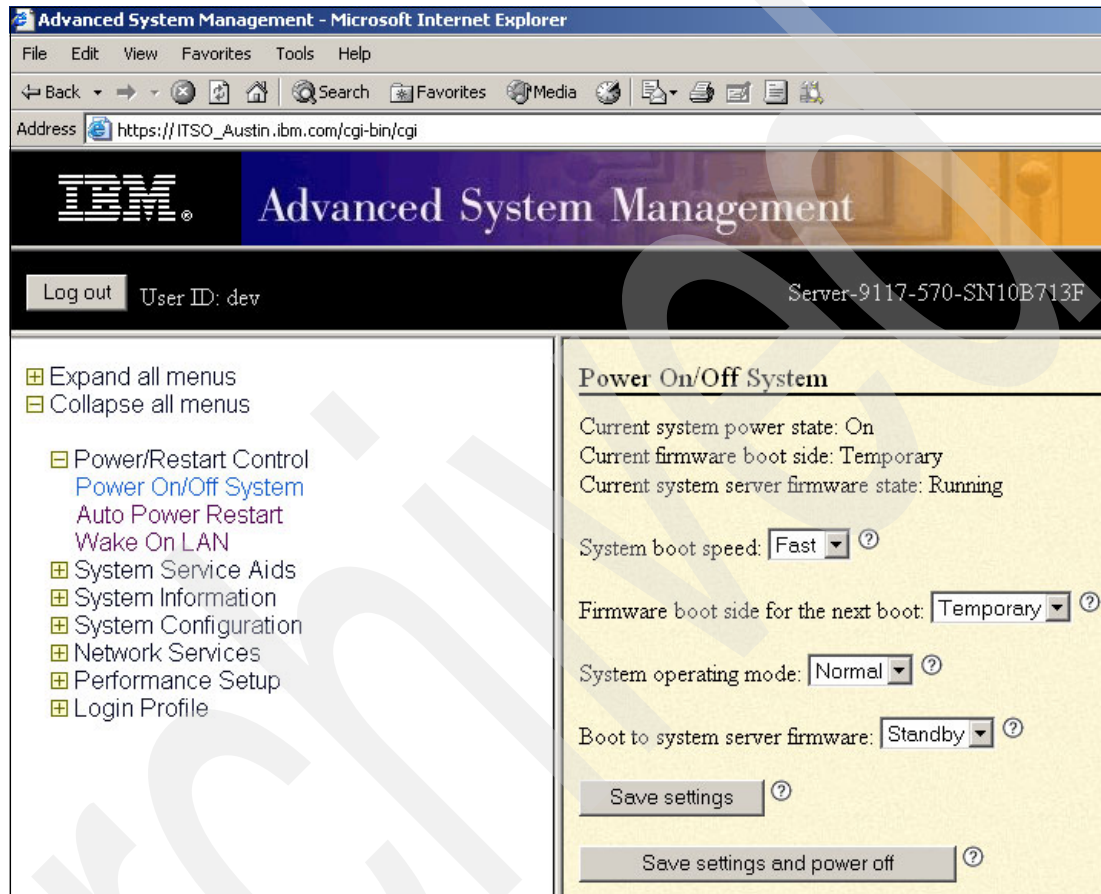


Figure 3-2 Advanced System Management main menu

3.2.2 Partition diagnostics

The diagnostics consist of stand-alone diagnostics, which are loaded from the DVD-ROM drive, and online diagnostics (available in AIX 5L).

- ▶ Online diagnostics, when installed, are resident with AIX 5L on the disk or server. They can be booted in single-user mode (service mode), run in maintenance mode, or run concurrently (concurrent mode) with other applications. They have access to the AIX 5L error log and the AIX 5L configuration data.
 - Service mode (requires service mode boot) enables the checking of system devices and features. Service mode provides the most complete checkout of the system resources. All system resources, except the SCSI adapter and the disk drives used for paging, can be tested.
 - Concurrent mode enables the normal system functions to continue while selected resources are being checked. Because the system is running in normal operation, some devices might require additional actions by the user or diagnostic application before testing can be done.

- Maintenance mode enables the checking of most system resources. Maintenance mode provides the exact same test coverage as Service Mode. The difference between the two modes is the way they are invoked. Maintenance mode requires that all activity on the operating system be stopped. The **shutdown -m** command is used to stop all activity on the operating system and put the operating system into maintenance mode.
- ▶ The System Management Services (SMS) error log is accessible from the SMS menu for tests performed through SMS programs. For results of service processor tests, access the error log from the service processor menu.

Note: Because the p5-560Q system has an optional DVD-ROM (FC 1994) and DVD-RAM (FC 1993), alternate methods for maintaining and servicing the system need to be available if the DVD-ROM or DVD-RAM is not ordered. It is possible to use Network Install Manager (NIM) server for this purpose.

3.2.3 Service Agent

Service Agent is an application program that operates on an IBM System p computer and monitors them for hardware errors. It reports detected errors, assuming they meet certain criteria for severity, to IBM for service with no intervention. It is an enhanced version of Service Director™ with a graphical user interface.

Key things you can accomplish using Service Agent for System p5, pSeries, and RS/6000 include:

- ▶ Automatic VPD collection
- ▶ Automatic problem analysis
- ▶ Problem-definable threshold levels for error reporting
- ▶ Automatic problem reporting; service calls placed to IBM without intervention
- ▶ Automatic client notification

In addition:

- ▶ Commonly viewed hardware errors. You can view hardware event logs for any monitored machine in the network from any Service Agent host user interface.
- ▶ High-availability cluster multiprocessing (HACMP) support for full fallback. Includes high-availability cluster workstation (HACWS) for 9076.
- ▶ Network environment support with minimum telephone lines for modems.
- ▶ Provides communication base for performance data collection and reporting tool Performance Management (PM/AIX). For more information about PM/AIX, see:

<http://www.ibm.com/servers/aix/pmaix.html>

Machines are defined by using the Service Agent user interface. After the machines are defined, they are registered with the IBM Service Agent Server (SAS). During the registration process, an electronic key is created that becomes part of your resident Service Agent program. This key is used each time the Service Agent places a call for service. The IBM Service Agent Server checks the current client service status from the IBM entitlement database; if this reveals that you are not on Warranty or MA, the service call is refused and posted back using an e-mail notification.

Service agent can be configured to connect to IBM either using modem or network connection. In any case, the communication is encrypted and strong authentication is used.

Service Agent sends outbound transmissions only and does not allow any inbound connection attempts. Only hardware machine configuration, machine status or error information is transmitted. Service Agent does not access or transmit any other data on the monitored systems.

Three principal ways of communication are possible:

- ▶ Dial-up using attached modem device (uses the AT&T Global Network dialer for modem access, does not accept incoming calls to modem)
- ▶ VPN (IPsec is used in this case)
- ▶ HTTPS (can be configured to work with firewalls and authenticating proxies)

The following Figure 3-3 show possible communication paths how can an IBM System p5 system be configured to utilize all features of Service Agent. The shown communication to IBM support can be either modem or network. If HMC is present, Service Agent is an integral part of it and if activated will collect hardware related information and error messages about the whole system and partitions. If software level information (like performance data for example) is also required, Service Agent can also be installed on any of the partitions and can be configured to act as either gateway and connection manager or a client. Gateway and connection manager gathers data from clients and communicates to IBM on behalf of them.

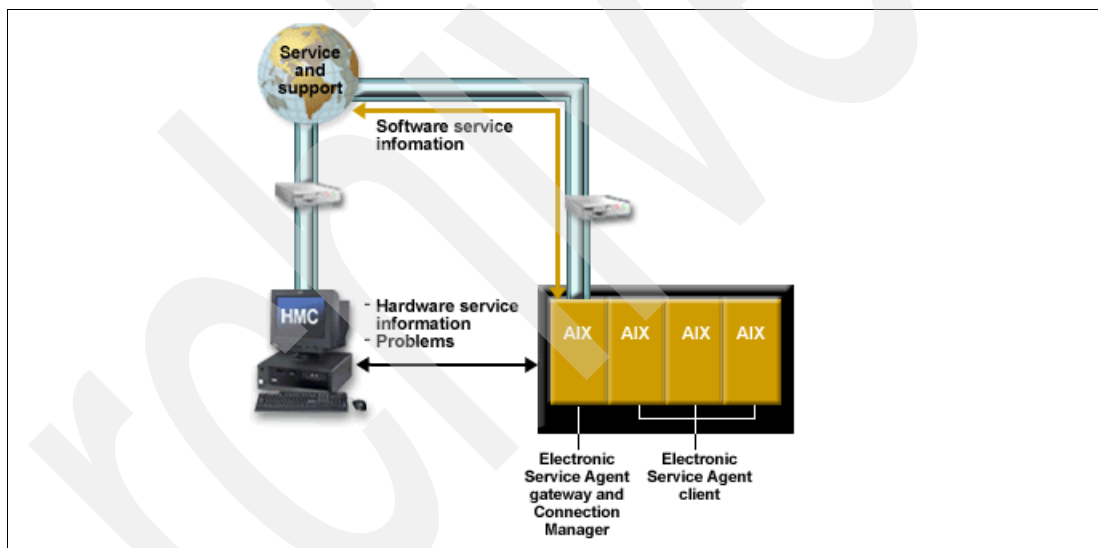


Figure 3-3 Service agent and possible connections to IBM

Additional services provided by Service Agent:

- ▶ My Systems: Client and IBM persons authorized by client can view HW information and error messages gathered by Service Agent on Electronic Services WWW pages (<http://www.ibm.com/support/electronic>)
- ▶ Premium Search: Search service using information gathered by Service Agents (paid service that requires special contract).
- ▶ Performance Management: Service Agent provides means for collecting long term performance data. The data are collected in reports accessed by client on WWW pages of Electronic Services (paid service that requires special contract).

You can download the latest version of Service Agent at:

ftp://ftp.software.ibm.com/aix/service_agent_code

Service Focal Point

Traditional service strategies become more complicated in a partitioned environment. Each logical partition reports errors it detects, without determining whether other logical partitions also detect and report the errors. For example, if one logical partition reports an error for a shared resource, such as a managed system power supply, other active logical partitions might report the same error. The Service Focal Point application helps you to avoid long lists of repetitive call-home information by recognizing that these are repeated errors and correlating them into one error.

Service Focal Point is an application on the HMC that enables you to diagnose and repair problems on the system. In addition, you can use Service Focal Point to initiate service functions on systems and logical partitions that are not associated with a particular problem. You can configure the HMC to use the Service Agent call-home feature to send IBM event information. Service Focal Point is available also in Integrated Virtualization Manager. It allows you to manage serviceable events, create serviceable events, manage dumps, and collect vital product data (VPD) but no reporting via Service Agent is possible.

3.2.4 IBM System p5 firmware maintenance

IBM System p5, pSeries, and RS/6000 Customer-Managed Microcode is a methodology that enables you to manage and install microcode updates on System p5, pSeries, and RS/6000 systems and associated I/O adapters. The IBM System p5 Microcode can be installed either from HMC or from a running partition, for update details, see 2.13.3, “System firmware” on page 60.

If you use an HMC to manage your server, you can use the HMC interface to view the levels of server firmware and power subsystem firmware that are installed on your server, and are available to download and install.

Each System p5 server has the following levels of server firmware and power subsystem firmware:

- ▶ **Installed level** – This is the level of server firmware or power subsystem firmware that has been installed and will be installed into memory after the managed system is powered off and powered on. It is installed on *i* side of system firmware, for additional discussion about firmware sides see 2.13.4, “Service processor” on page 63.
- ▶ **Activated level** – This is the level of server firmware or power subsystem firmware that is active and running in memory.
- ▶ **Accepted level** – This is the backup level of server or power subsystem firmware. You can return to this level of server or power subsystem firmware if you decide to remove the installed level. It is installed on *p* side of system firmware, for additional discussion about firmware sides see 3.2.1, “Service processor” on page 73.

IBM introduced the Concurrent Firmware Maintenance (CFM) function on System p5 servers in system firmware level 01SF230_126_120, which was released on 16 June 2005. This function supports non disruptive system firmware service packs to be applied to the system concurrently (without requiring a reboot to activate changes). For systems that are not managed by an HMC, the installation of system firmware is always disruptive.

The concurrent levels of system firmware might, on occasion, contain fixes that are known as deferred. These deferred fixes can be installed concurrently, but will not be activated until the next IPL. Deferred fixes, if any, will be identified in the Firmware Update Descriptions table of this document. For deferred fixes within a service pack, only the fixes in the service pack which cannot be concurrently activated are deferred.

Use the following information as a reference to determine whether your installation will be concurrent or disruptive.

Figure 3-4 shows the system firmware file naming convention.




Figure 3-4 System firmware file naming convention

An installation is disruptive if:

- ▶ The release levels (XXX) of currently installed and new firmware are different.
- ▶ The service pack level (YYY) and the last disruptive service pack level (ZZZ) are equal in new firmware.

Otherwise, an installation is concurrent if:

- ▶ If the service pack level (YYY) of new the firmware is higher than the service pack level currently installed on the system and the above conditions for disruptive installation are not met.

3.3 Cluster solution

Today's IT infrastructure requires that servers meet increasing demands, while offering the flexibility and manageability to rapidly develop and deploy new services. IBM clustering hardware and software provide the building blocks, with availability, scalability, security, and single-point-of-management control, to satisfy these needs. The advantages of clusters are:

- ▶ Large-capacity data and transaction volumes, including support of mixed workloads
- ▶ Scale-up (add processors) or scale-out (add servers) without downtime
- ▶ Single point-of-control for distributed and clustered server management
- ▶ Simplified use of IT resources
- ▶ Designed for 24 x 7 access to data applications
- ▶ Business continuity in the event of disaster

The POWER5+ processor-based AIX 5L and Linux cluster targets scientific and technical computing, large-scale databases, and workload consolidation. IBM Cluster Systems Management software (CSM) is designed to provide a robust, powerful, and centralized way to manage a large number of POWER5 processor-based servers, all from one single point-of-control. Cluster Systems Management can help lower the overall cost of IT ownership by helping to simplify the tasks of installing, operating, and maintaining clusters of servers. Cluster Systems Management can provide one consistent interface for managing both AIX 5L and Linux nodes (physical systems or logical partitions), with capabilities for remote parallel network install, remote hardware control, and distributed command execution.

Cluster Systems Management for AIX 5L and Linux on POWER processor-based servers is supported on the p5-560Q. For hardware control, an HMC is required. One HMC can also control several servers that are part of the cluster. If a p5-560Q that is configured in partition

mode (with physical or virtual resources) is part of the cluster, all partitions must be part of the cluster.

Monitoring is much easier to use, and the system administrator can monitor all of the network interfaces, not just the switch and administrative interfaces. The management server pushes information out to the nodes, which releases the management server from having to trust the node. In addition, the nodes do not have to be network-connected to each other. This means that giving root access on one node does not mean giving root access on all nodes. The base security setup is all done automatically at install time.

For information regarding the IBM Cluster Systems Management for AIX 5L, HMC control, cluster building block servers, and cluster software available, visit the following links:

► Cluster 1600

<http://www.ibm.com/servers/eserver/clusters/hardware/1600.html>

► Cluster 1350™

<http://www.ibm.com/servers/eserver/clusters/hardware/1350.html>

The CSM ships with AIX 5L itself (a 60-day Try and Buy license is shipped with AIX). The CSM client side is automatically installed and ready when you install AIX, so each system or logical partition is cluster-ready.

The CSM V1.4 on AIX 5L and Linux introduces an optional IBM CSM High Availability Management Server (HA MS) feature, which is designed to allow automated failover of the CSM management server to a backup management server. In addition, sample scripts for setting up NTP¹, and network tuning (AIX 5L only) configurations, and the capability to copy files across nodes or node groups in the cluster can improve cluster ease of use and site customization.

¹ Network Time Protocol

Archived

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 83. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *Advanced POWER Virtualization on IBM eServer p5 Servers: Architecture and Performance Considerations*, SG24-5768
- ▶ *Advanced POWER Virtualization on IBM System p5*, SG24-7940
- ▶ *IBM eServer p5 510 Technical Overview and Introduction*, REDP-4001
- ▶ *IBM eServer p5 520 Technical Overview and Introduction*, REDP-9111
- ▶ *IBM eServer p5 550 Technical Overview and Introduction*, REDP-9113
- ▶ *IBM eServer p5 570 Technical Overview and Introduction*, REDP-9117
- ▶ *IBM eServer p5 590 and 595 System Handbook*, SG24-9119
- ▶ *IBM eServer p5 590 and 595 Technical Overview and Introduction*, REDP-4024
- ▶ *IBM eServer pSeries Sizing and Capacity Planning: A Practical Guide*, SG24-7071
- ▶ *IBM System p5 505 Express Technical Overview and Introduction*, REDP-4079
- ▶ *IBM System p5 510 and 510Q Technical Overview and Introduction*, REDP-4136
- ▶ *IBM System p5 520 and 520Q Technical Overview and Introduction*, REDP-4137
- ▶ *IBM System p5 560Q Technical Overview and Introduction*, REDP-4139
- ▶ *LPAR Simplification Handbook*, SG24-7231
- ▶ *Partitioning Implementations for IBM eServer p5 Servers*, SG24-7039
- ▶ *Problem Solving and Troubleshooting in AIX 5L*, SG24-5496
- ▶ *Virtual I/O Server Integrated Virtualization Manager*, SG24-4061

Other publications

These publications are also relevant as further information sources:

- ▶ *7014 Series Model T00 and T42 Rack Installation and Service Guide*, SA38-0577, contains information regarding the 7014 Model T00 and T42 Rack, in which this server can be installed.
- ▶ *IBM eServer Hardware Management Console for pSeries Installation and Operations Guide*, SA38-0590, provides information to operators and system administrators on how to use an IBM Hardware Management Console for pSeries (HMC) to manage a system. It also discusses the issues associated with logical partitioning planning and implementation.
- ▶ *IBM System p5, IBM eServer p5 and i5, and OpenPower Planning*, SA38-0508

- ▶ *Planning for Partitioned-System Operations*, SA38-0626, provides information to planners, system administrators, and operators about how to plan for installing and using a partitioned server. It also discusses some issues associated with the planning and implementing of partitioning.
- ▶ *RS/6000 and eServer pSeries Diagnostics Information for Multiple Bus Systems*, SA38-0509, contains diagnostic information, service request numbers (SRNs), and failing function codes (FFCs).
- ▶ *System p5, eServer p5 Customer Service Support and Troubleshooting*, SA38-0538, contains information regarding slot restrictions for adapters that can be used in this system.
- ▶ *System Unit Safety Information*, SA23-2652, contains translations of safety information used throughout the system documentation.

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ AIX 5L operating system maintenance packages downloads
<http://www.ibm.com/servers/eserver/support/unixservers/aixfixes.html>
- ▶ News on computer technology
<http://www.ibm.com/chips/micronews>
- ▶ Copper circuitry
<http://domino.research.ibm.com/comm/pr.nsf/pages/rsc.copper.html>
- ▶ IBM Systems Hardware Information Center documentation
<http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp>
- ▶ IBM Systems Information Centers
<http://publib.boulder.ibm.com/eserver/>
- ▶ IBM microcode downloads
<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>
- ▶ Support for IBM System p servers
<http://www.ibm.com/servers/eserver/support/unixservers/index.html>
- ▶ Technical help database for AIX 5L
<http://www14.software.ibm.com/webapp/set2/srchBroker/views/srchBroker.jsp?rs=111>
- ▶ IBMlink
<http://www.ibm.link.ibm.com>
- ▶ Linux for IBM System p5
<http://www.ibm.com/systems/p/linux/>
- ▶ Microcode Discovery Service
<http://www14.software.ibm.com/webapp/set2/mds/fetch?page=mds.html>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Archived



IBM System p5 560Q

Technical Overview and Introduction



Redpaper

Finer system granulation using Micro-Partitioning technology to help lower TCO

Modular midrange solution for managing on demand business

New POWER5+ processor options using DDR2 memory technology

This document is a comprehensive guide that covers the IBM System p5 560Q UNIX servers. It introduces major hardware offerings and discusses their prominent functions.

Professionals wishing to acquire a better understanding of IBM System p products should read this document. The intended audience includes:

- Clients
- Sales and marketing professionals
- Technical support professionals
- IBM Business Partners
- Independent software vendors

This document expands the current set of IBM System p documentation by providing a desktop reference that offers a detailed technical description of the p5-560Q system.

This publication does not replace the latest IBM System p marketing materials and tools. It is intended as an additional source of information that you can use, together with existing sources, to enhance your knowledge of IBM server and workstation solutions.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks