Performance of the IBM System p5 520, 550 and 550Q

October 4, 2005

Duc J. Vianney

Jacob Thomas

William A. Maron

IBM Systems and Technology Group

Abstract

The IBM System p5™ 520, 550 and 550Q are the latest IBM POWER5™ processor-based servers designed for commercial and high performance compute-intensive workloads. The p5-520 and p5-550 servers are 2-way and 4-way Dual-Chip Module systems running at 1.9 GHz while the p5-550Q is 4-way and 8-way 1.5 GHz Quad-Core Module systems. All three systems support the AIX 5L™ (AIX 5L V5.3 and AIX 5L V5.2) and Linux® (SUSE LINUX Enterprise Server 9 (SLES 9) for POWER™, and Red Hat Enterprise Linux AS 4 (RHEL AS 4) for POWER operating systems. In this paper, we examine key features of the three systems and discuss their performance under such industry standard benchmarks as SPEC CPU2000, SPEC OMP2001, SPECjbb2000, SPECjbb2005 and LINPACK. The measured benchmark results demonstrate the leading edge performance of the new IBM POWER5+™ systems in commercial and high performance workloads. Following is the content of this paper.

- Introduction
- POWER5 Technology
- New System Features
- Selected Industry Standard Benchmarks
- Measured Industry Standard Benchmark Results
- Summary
- Appendix A - IBM System p5 520 Features
- Appendix B – IBM System p5 550 and p5 550Q Features
- Appendix C – Benchmark Descriptions

Introduction

The IBM System p5 520 and 550 are entry-level servers based on the IBM POWER5+ 64-bit Dual-Chip Module (DCM). The DCM forms the basic building block for low-end POWER5 servers. The DCM design consists of one dual-core chip with 1.9MB on-chip L2 cache, an integrated memory controller, a 36MB off-chip L3 cache, and the high performance switch logic that connect to other modules. The L2 cache and L3 cache are shared by the dual cores in the DCM. The 4-way p5-550 differs from the 2-way p5-520 only in configuration such as the number of DCMs / processors (two vs. four), the maximum amount of memory (64GB vs. 32GB), and the maximum number of supported I/O drawers (eight vs. four). Both systems can be packaged as deskside system unit for used in stand-alone or 4U rack mountable.

520 Deskside

The p5-550Q is also an entry-level server based on the 64-bit POWER5+ Quad-Core Module (QCM). The p5-550Q is offered in one and two QCMs. Each QCM consists of two dual-core chips with two 1.9MB on-chip L2 caches, an integrated memory controller and two 36MB off-chip L3 caches.

The POWER5+ processor technology provides faster clock and architecture enhancements, a new faster Quad-Core Module for higher n-way density, Double Data Rate memory (DDR2), and a faster I/O subsystem. These features will be discussed in a later section.

POWER5 Technology

The multi-chip module design POWER5 processors incorporate several new technologies designed to improve performance and scalability of server workloads. These include simultaneous multithreading, Micro-Partitioning™, additional renaming registers, improved cache structure with Level 3 (L3) cache on module and L3 directory and controller on chip, large Level 2 (L2) and L3 cache sizes with lower latency, enhanced memory loads and stores with on-chip memory controller with lower latency to main memory, faster I/O buses, and dynamic power management [1,2]. Other improvements designed to support the POWER5 Non-Uniform Memory Access (NUMA) architecture are hardware instruction cache coherence, hardware synchronization of translation-look-aside buffer (TLB) invalidation operations, and hardware barrier synchronization register [3].

Simultaneous multithreading provides support for the execution of two instruction streams, or threads, on a single processor core. In the POWER5 design, all threads are simultaneously active, and there is no thread switch even when one thread has a cache miss. As a result, there is no thread switch overhead and all processing functional units are designed to be used to the fullest extent [2].

Micro-Partitioning technology introduces the concept of shared processor partition by providing the ability to share a single processor core between multiple partitions as opposed to the dedicated partitions where no core sharing is allowed [4].

Register-renaming resources have been increased to support the additional register requirements of the second thread running in simultaneous multithreading mode. In the POWER5 processor, there are 120 general-purpose registers (GPRs) and 120 floating-point registers (FPRs) that form a shared register pool that can be renamed (mapped) to one of 32 GPRs or FPRs on a per-thread basis.

Figure 1 - POWER5 Chip Layout

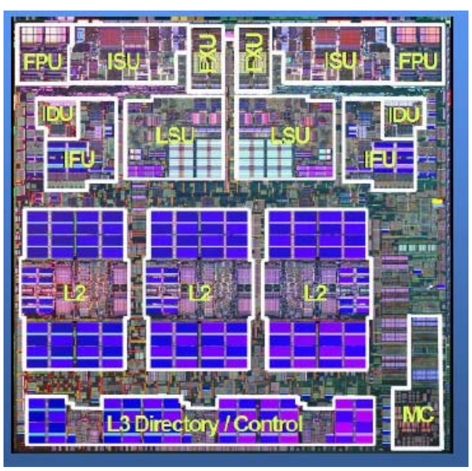

The POWER5 processor chip shown in Figure 1 contains two microprocessor cores, chip and system pervasive functions, core interface logic, a 1.9MB L2 cache and controls, a 36MB L3 cache directory and controls, and the fabric controller that controls the flow of information and control data between the L2 and L3 and between chips. The L3 is a victim cache of the L2 cache. Each core contains a 64 KB Level 1 (L1) instruction cache (I-cache), a 32 KB L1 data cache (D-cache), two fixed-point execution pipelines in the FXU unit, two floating-point execution pipelines in the FPU unit, two load/store execution pipelines in the LSU unit, an instruction fetch unit (IFU), an instruction decode unit (IDU), and an instruction sequencing unit (ISU). In addition not shown is a branch execution pipeline in the BRU unit and a condition register (CR) logical pipeline to perform logical operations on the condition in the CRU unit. See more details in [1,2,5].

The 64-bit POWER5 core is a speculative, out-of-order execution core coupled with a multilevel storage hierarchy. An overview of the POWER5 pipeline organization is shown in Figure 2.

Figure 2 – POWER5 Pipeline Structure

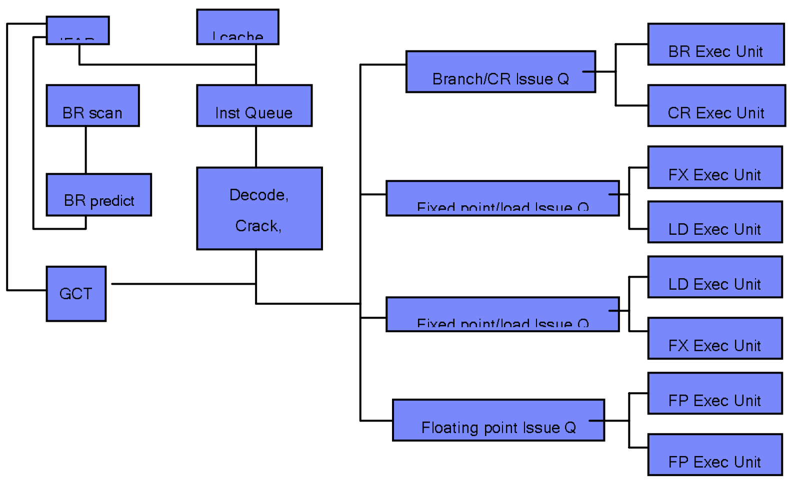

The pipeline structure comprises a master pipeline and several execution unit pipelines, all can progress independently from each other. As discussed in [9], the master pipeline presents speculative and in-order instructions to the mapping, sequencing and dispatch function, and is designed to ensure an orderly completion of the real execution path. It throws away any potential speculative results associated with miss predicted branches. The execution unit pipelines allow out-of-order issuing of speculative and non-speculative instructions.

Instructions dispatched in program order in groups are issued out of program order to the execution units, with a bias towards oldest operations first. Groups can consist of up to five instructions, and are always terminated by a branch instruction. In general, up to eight instructions are fetched from the I-cache and fed into the pipeline to go through various pipeline stages until completion. In each cycle, up to five instructions are pulled from the instruction fetch buffer Inst Queue and sent through the 3-stage instruction decode pipeline to form a dispatch group. Complex instructions are cracked into *internal operations* (IOPs) to allow for simpler inner core dataflow.

From dispatch to completion, instructions are tracked in group. Each group has an entry in the Global Completion Table (GCT). A group can contain up to 5 IOPs. During dispatch, various machine resources are assigned to instructions. These resources are:

- GCT: one entry per group
- FXU Issue queue: one entry per fixed-point or load/store IOP
- FPU Issue queue: one entry per floating-point IOP
- BRU Issue queue: one entry per branch IOP
- CRU Issue queue: one entry per CR IOP
- mappers: one entry per new destination (either GPR, FPR, XER, CR, Link/CNT)
- LRQ: one entry per load IOP
- SRQ: one entry per store IOP

LRQ is the Load Reorder Queue, a 32-entry queue which holds real addresses and tracks the order of loads. SRQ is the Store Reorder Queue, a 32-entry queue which tracks all stores active in the LSU. In Figure 2, IFAR represents the Instruction Fetch Address Register which handles program-counter functionality. The following machine resources are released at various stages:

- GCT: when the group completes, i.e., when all instructions in the oldest group finished their execution
- Issue queue: after the instruction have been successfully issued
- Mappers: when the group completes
- LRQ: when the group completes
- SRQ: when the store is sent to the storage subsystem

The branch prediction logic BR Scan scans all the fetched instructions looking for up to two branches per cycle. Depending upon the branch type found, various branch prediction mechanisms in Branch BR engage to help predict the target address of the branch or the branch direction or both.

New System Features

POWER5+ Technology

The POWER5+ processor technology provides several changes to the core organization including faster clock and architecture enhancements and a new faster QCM for higher n-way density. The QCM used in the p5-550Q consists of two POWER5+ chips; each is a dual-core for a total of four cores and two L3 chips in a module. This layout enables the System p5 entry servers to scale up to 8-way in a 4U package.

The memory for POWER5+ is offered as pluggable DIMMs. Only 533 MHz DDR2 memory DIMMs are supported on p5-520, p5-550 and p5-550Q. A brief description of the memory subsystem is as follows.

- Integrated memory controller, 2 x 8 bytes READ and 2 x 2 bytes WRITE bus
- Eight DIMM slots per processor card, increments in pair or quad
- Type: DDR2 533 MHz.
- Min-Max DDR2: 1GB - 64GB
- Package: DDR2: 276-pin
- Enhanced memory RAS features including ECC, bit-steering, Chipkill™ detection and correction

For a detailed specification on p5-520, p5-550 and p5-550Q, see Appendix A and Appendix B.

IBM has developed a new technology colloquially known as “Chipkill Protect ECC DIMMs” which will allow an entire DRAM chip on a DIMM to fail while the system continues to function. These new DIMMs have been carefully designed so that there is no performance degradation over single-error correct (SEC) or standard ECC DIMMs.

Data transfers to and from an SDRAM DIMM use a synchronous clock signal to establish timing. The Double Data Rate (DDR) memory techniques increase the data rate by transferring data on both the rising edge and the falling edge of the clock signal. DDR DIMMs use a “2x” prefetch scheme so that two sets of 64-bit data are referenced simultaneously. Logic on the DIMM multiplexes the two 64-bit results (plus ECC bits) to appear on each of the rising and falling edges of the clock signal. Thus, two data transfers can be performed during one clock period.

DDR2 is the new generation of DDR technology. The primary benefit is the potential for faster throughput. Currently, DDR2 operates at data transfer rates starting at 400 MHz (the upper limit for DDR) and 533 MHz. In addition, the DDR2 improves the power consumption of the DIMM because it works on a lower voltage. DDR operates at a range of 2.5v to 2.8v, whereas DDR2 only requires 1.8v. DDR2 consumes less power than DDR and offers a higher range of throughput because it has halved the speed of the memory core (thereby reducing power consumption) but has offset that by doubling the number of prefetches from the memory core to the I/O buffers (from two to four). [10] For a detailed discussion of memory performance and DDR memory, see [10]. The 400 MHz DDR2 is not available with the current release.

The I/O subsystems include the P5IOC-based I/O subsystem which offers higher bandwidth and lower latency, and support for new industry standard IO subsystems like InfiniBand® (available November 18, 2005) and RIOG.

- Eight hot-swappable DASD
- Three media bays: two slimline bays and one half-high bay
- Fibe hot-plug PCI-X slots with Enhanced Error Handling (EEH). PCI-X adapters installed in slot 4 and slot 5 must be short card.
- Redundant and hot-plug power with dual line cord (optional)
- Redundant and hot-plug cooling
- Dynamic LPAR (require optional HMC console)
- 2 GX+ slots for external connectivity or I/O drawer expansion (Two GX+ slots and two PCI-X slots share the same physical locations)
- Integrated Dual Channel SCSI Ultra320 controller
- Dual Channel Gigabit Ethernet controller
- Flexible Service Processor (FSP) for enhanced reliability and remote system management
- Deskside or rack drawer package

Selected Industry Standard Benchmarks

The following industry standard benchmarks were used to evaluate the performance of the p5-520, p5-550 and p5-550Q servers: SPEC CPU2000, SPEC OMP2001, SPECjbb2000, SPECjbb2005 and LINPACK. The benchmark descriptions are detailed in [6], and summarized in Appendix C. For more detailed information on the benchmark and published results, visit the benchmark Web site [7,8].

- SPEC CPU2000: Compute intensive, integer and floating-point performance
- SPEC OMP2001: Compute-intensive parallel workloads
- SPECjbb2000 / SPECjbb2005: Server-side Java™ workloads
- LINPACK: Solving a dense system of linear equations

Performance Results

p5-520, p5-550 and p5-550Q Compared Against Best of Published Results

Table 1 shows the performance of those three new systems against the best of published results (BoPR) systems whose results are available publicly [7,8]. Detailed configurations are also available at the corresponding benchmark Web site. Based on the results available as of October 4, 2005, we observe the following:

- The p5-520 produces the highest score in SPECfp2000.
- The p5-520 is the best 2-way performer in SPECfp_rate2000, LINPACK HPC, SPECompMpeak2001, and SPECjbb2005.
- The p5-550 holds three best 4-way positions in SPEC_int_rate2000, SPECfp_rate2000, LINPACK HPC, SPECjbb2000 and SPECjbb2005.
- The p5-550Q holds the best 8-way position in SPECjbb2005.

Table 1 - p5-520, p5-550 and p5-550Q vs. Best of Published Results Systems

Benchmarks	Cores	p5-520 1.9GHz	BoPR Scores	BoPR System Name	BoPR Processors
SPECint2000	1-way	1,513	1,970	ASUS A8N-SLI Deluxe, AMD Athlon (TM) 64 FX-57	2.8 GHz AMD Athlon (1)
SPECfp2000	1-way	3,030	2,801	HITACHI BladeSymphony (1.66 GHz / 9MB Itanium 2)	1.66 GHz Intel® Itanium® 2 (2)
SPECint_rate2000	2-way	39.6	40.8	Altos G5350 (AMD Opteron (TM) 252)	2.6 GHz AMD Opteron™ (3)
SPECfp_rate2000	2-way	67.6	56.7	HITACHI BladeSymphony (1.66 GHz / 9MB Itanium 2)	1.66 GHz Intel Itanium 2 (4)
Linpack HPC	2-way	14,310	13,270	IBM @server® BladeCenter JS20 (2.2GHz PowerPC®)	2.2 GHz, IBM JS20 (5)
SPECompMpeak2001	2-way	8,174	7,612	Sun Fire X4100	2.8 GHz Sun X4100 (6)
SPECjbb2000	2-way	99,844	105,296	PowerEdge SC1425	3.6 GHz Xeon™ DP (7)
SPECjbb2005	2-way	32,820	27,004	Sun Fire X4100	2.8 GHz Sun X4100 (8)

Benchmarks	Cores	p5-550 1.9GHz	BoPR Scores	BoPR Systems	BoPR Processors
SPECint_rate2000	4-way	78.5	78.4	ProLiant BL45p (AMD Opteron (TM) 852)	2.6 GHz AMD Opteron (9)
SPECfp_rate2000	4-way	133	130	IBM @server p5 570 (1.9 GHz, 4 CPU)	1.9 GHz IBM p5-570 (10)
LINPACK HPC	4-way	28,490	27,520	IBM @server p5 570 (1.9 GHz POWER5)	1.9 GHz IBM p5-570 (11)
SPECompMpeak2001	4-way	15,392	16,096	IBM @server p5 570 (1900 MHz, 4 CPU)	1.9 GHz IBM p5-570 (12)
SPECjbb2000	4-way	190,445	181,496	Devices Tyan Computer Corporation S4882-D	2.6 GHz AMD Opteron (13)
SPECjbb2005	4-way	61,789	37,034	Sun Fire V40z	2.6 GHz AMD Opteron (14)

Benchmarks	Cores	p5-550Q 1.5GHz	BoPR Scores	BoPR Systems	BoPR Processors
SPECint_rate2000	8-way	124	167	IBM @server p5 575 (1.9 GHz, 8 CPU)	1.9 GHz IBM p5-575 (15)
SPECfp_rate2000	8-way	178	282	IBM @server p5 575 (1.9 GHz, 8 CPU)	1.9 GHz IBM p5-575 (16)
LINPACK HPC	8-way	44,680	56,780	IBM @server p5 575 (1.9 GHz POWER5)	1.9 GHz IBM p5-575 (17)
SPECompMpeak2001	8-way	20,122	28,035	IBM @server p5 575 (1.9 GHz, 8 CPU)	1.9 GHz IBM p5-575 (18)
SPECjbb2000	8-way	294,315	328,996	IBM @serverr p5 570 (1.9 GHz)	1.9 GHz IBM p5-570 (19)
SPECjbb2005	8-way	91,806	N/A	N/A	N/A

SPECsfs97_R1.v3	8-way	118,391	145,362	IBM @server p5 570 (1.9 GHz, TCP)	1.9 GHz IBM p5-570 (20)

Sources:

(1) http://www.spec.org/cpu2000/results/res2005q2/cpu2000-20050613-04264.html
(2) http://www.spec.org/cpu2000/results/res2005q3/cpu2000-20050628-04342.html
(3) http://www.spec.org/cpu2000/results/res2005q3/cpu2000-20050724-04388.html
(4) http://www.spec.org/cpu2000/results/res2005q3/cpu2000-20050628-04343.html
(5) http://www.netlib.org/benchmark/performance.pdf
(6) http://www.spec.org/omp/results/res2005q3/omp2001-20050909-00198.html
(7) http://www.spec.org/jbb2000/results/res2005q1/jbb2000-20050315-00312.html
(8) http://www.spec.org/jbb2005/results/res2005q3/jbb2005-20050830-00013.html
(9) http://www.spec.org/cpu2000/results/res2005q2/cpu2000-20050502-04060.html
(10) http://www.spec.org/cpu2000/results/res2004q3/cpu2000-20040712-03230.html
(11) http://www.netlib.org/benchmark/performance.pdf
(12) http://www.spec.org/omp/results/res2004q4/omp2001-20041014-00168.html
(13) http://www.spec.org/jbb2000/results/res2005q3/jbb2000-20050705-00347.html
(14) http://www.spec.org/jbb2005/results/res2005q2/jbb2005-20050613-00003.html
(15) http://www.spec.org/cpu2000/results/res2005q1/cpu2000-20050207-03765.html
(16) http://www.spec.org/cpu2000/results/res2005q1/cpu2000-20050207-03764.html
(17) http://www.netlib.org/benchmark/performance.pdf
(18) http://www.spec.org/omp/results/res2005q1/omp2001-20050207-00184.html
(19) http://www.spec.org/jbb2000/results/res2004q3/jbb2000-20040712-00247.html
(20) http://www.spec.org/sfs97r1/results/res2004q3/sfs97r1-20040712-00191.html

NOTES:

- Table 1: Current as of October 4, 2005. Sources: www.spec.org and http://www.netlib.org/benchmark/performance.pdf
- Disclaimers: All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying.

Summary

The POWER5+ processor continues the performance legacy of the POWER5 processor with faster clock, architecture enhancements, faster memory and innovative QCM design allowing the System p5 entry servers to scale up to 8-way leveraging the simultaneous multithreading efficiency in the POWER5 processor. The new QCM configuration provides more processing threads which will benefit such applications like databases, OLTP transactions, Web-commerce transaction platform and Internet applications based on Java. In addition, the virtualization technology along with the scalability of Power Architecture™ demonstrate the capability of the new entry-level POWER5+ servers as a low costs consolidation platform for multiple applications and operating systems.

References

[1] Sinharoy, B. et al., POWER5 system microarchitecture, *IBM JRD 49-4/5, July/September* 2005.

Available: http://www.research.ibm.com/journal/rd/494/sinharoy.html

[2] Mathis, H. et al., Characterization of simultaneous multithreading efficiency in POWER5, *IBM JRD 49-4/5, July/September* 2005. Available: http://www.research.ibm.com/journal/rd/494/mathis.html

[3] Mackeras P. et al, Operating system exploitation of the POWER5 system, *IBM JRD 49-4/5*, May 2005. Available: http://www.research.ibm.com/journal/rd/494/mackerras.html

[4] Armstrong, W. et al., Advanced virtualization capabilities of POWER5 systems, *JRD 49-4/5, July/September* 2005. Available; http://www.research.ibm.com/journal/rd/494/armstrong.html

[5] C. Grassl, New IBM components for HPCx, *HPCx Annual Seminar*, December 2003. Available: http://www.hpcx.ac.uk/about/events/annual2003/Grassl.pdf

[6] Understanding IBM @server xSeries® benchmarks, *IBM REDP-3957*, May 2005.

[7] SPEC Benchmarks. Available: http://www.spec.org/

[8] LINPACK HPC Benchmark. Available: http://www.netlib.org/LINPACK/

[9] O. Herescu, B. Olszewski, and H. Hua, Performance Workloads Characterization on POWER5 with Simultaneous Multi Threading Support, *Eighth Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW-8)* February 2005. Available; http://tesla.hpl.hp.com/caecw05

[10] Tuning IBM @server xSeries servers for performance, *IBM Redbook SG-24-5287*, December 2004.

Appendix 1 - IBM System p5 520 Features

IBM System p5 520 (9131-52A)

- Deskside and 4U rack-mount server
- Based on IBM @server p5 520 system design
- P5-IOC planar
- Processor options
- 2-way 1.9 GHz POWER5+ DCM (36MB L3)
- 1GB - 32GB DDR2 memory
- 64 DIMM slots per node
- Six hot-plug PCI-X slots (one is DDR 266 MHz)
- One GX slot
- Four standard + four optional hot-swappable 3.5-inch disk drive bays
- Three media bays
- Other functions:
- Integrated service processor
- Dual-ported Ethernet 10/100/1000 Mbps controller
- Two HMC ports
- Two USB ports
- Dual-ported integrated Ultra320 SCSI controller (internal) with optional RAID Enablement Card
- Two System ports
- Two SPCN ports
- Redundant cooling
- Optional redundant power
- Optional Advanced POWER Virtualization feature
- Up to four 7311-D20 I/O drawers
- With RIO-2 adapter (no integrated RIO-2 ports)
- Express offerings available
- p5-520 Express, AIX 5L Edition - include a zero priced processor entitlement and lower priced AIX 5L license(s)
- p5-520 Express, OpenPower™ Edition - include a zero priced processor entitlement and lower priced SLES 9 or RHEL AS 4 subscription
- AIX 5L V5.2 and V5.3, SLES 9, RHEL AS 3, RHEL AS 4
- Planned GA October 14, 2005, except:
- Planned GA November 18, 2005 for:
- PCI-X Dual-port 4x InfiniBand Adapter (FC 1820)
- 3 meter 4x InfiniBand Cable (FC 1835)
- 8 meter 4x InfiniBand Cable (FC 1836)
- 1.5 meter 4x InfiniBand Cable (FC 1839)
- Planned GA March 31, 2006 - 8GB (2 x 4GB) DDR2 memory (FC 1934)

Appendix 2 - IBM System p5 550 and 550Q Features

IBM System p5 550 and 550Q (9133-55A)

- Deskside and 4U rack-mount server
- Based on IBM @server p5 550 system design
- P5-IOC planar
- Processor options
- 2- or 4-way 1.9 GHz POWER5+ DCM (36MB L3 per processor card) (p5-550)
- 4- or 8-way 1.5 GHz POWER5+ QCM (72MB L3 per processor card) (p5-550Q)
- 1GB - 64GB DDR2 memory
- 64 DIMM slots per node
- Five hot-plug PCI-X slots (one is DDR 266 MHz)
- Two GX slots occupy same space as the two short PCI-X slots
- Four standard + four optional hot-swappable 3.5-inch disk drive bays
- Three media bays
- Other functions:
- Integrated service processor
- Dual-ported Ethernet 10/100/1000 Mbps controllers
- Two HMC ports
- Two USB ports
- Dual-ported integrated Ultra320 SCSI controller (internal) with optional RAID Enablement Card
- Two System ports
- Two SPCN ports
- Redundant cooling
- Optional redundant power
- Optional Advanced POWER Virtualization feature
- Up to eight 7311-D20 I/O drawers
- With RIO-2 adapters (no integrated RIO-2 ports)
- Express offerings available
- p5-550 Express and p5-550Q Express, AIX 5L Edition - include one, two or four zero priced processor entitlements and lower priced AIX 5L license(s)
- p5-520 Express and p5-550Q Express, OpenPower Edition - include one, two or four zero priced processor entitlements and lower priced SLES 9 or RHEL AS 4 subscription
- AIX 5L V5.2 and V5.3, SLES 9, RHEL AS 3, RHEL AS 4
- Planned GA October 14, 2005, except:
- Planned GA November 18, 2005 for:
- PCI-X Dual-port 4x InfiniBand Adapter (FC 1820)
- 3 meter 4x InfiniBand Cable (FC 1835)
- 8 meter 4x InfiniBand Cable (FC 1836)
- 1.5 meter 4x InfiniBand Cable (FC 1839)
- Planned GA March 31, 2006 - 8GB (2 x 4GB) DDR2 memory (FC 1934)

Appendix C - Benchmark Descriptions

SPEC CPU2000 provides a measure of compute-intensive performance across a wide range of hardware. It measures system speed and throughput for single-processor, symmetric-multiprocessor, and cluster systems. SPEC CPU2000 contains two sets (or suites) of benchmarks: CINT2000 for measuring compute-intensive integer performance and CFP2000 for compute-intensive floating point performance.

CINT2000 provides two performance metrics. The throughput metric, SPECint_rate2000, measures the number of tasks a computer can complete within a given time interval. The metric is the geometric mean of the 12 SPEC rates from the SPEC integer tests (CINT2000). SPECint_base_rate2000 is the result of the same tests as CINT2000 with a maximum of four compiler options that must be used in all 12 tests.

The speed metric, SPECint2000, measures how fast a machine completes the running of the CINT2000 suite. The speed run is generally being done on a single CPU system. The result is the geometric mean of 12 tests that comprise the CINT2000 benchmark suite. All of these are written in C language except for one which is in C++. SPECint_base2000 is the result of the same tests as CINT2000 with a maximum of four compiler options that must be used in all 12 tests.

CFP2000 also provides two performance metrics. The throughput metric, SPECfp_rate2000, measures the number of tasks a computer can complete within a given time interval. The result is the geometric mean of 14 tests, all written in FORTRAN and C languages that are included in the CFP2000 benchmark suite. SPEC_base_rate2000 is the result of the same tests as CFP2000 with a maximum of four compiler options that must be used in all 14 tests.

The speed metric, SPECfp2000, measures how fast a machine completes the running of the CFP2000 suite. SPECfp_base2000 is the result of the same tests as CFP2000 with a maximum of four compiler options that must be used in all 14 tests.

SPEC OMP2001 is also a CPU intensive benchmark suite comprising workloads that uses Open MP standard. It measures compute intensive parallel workloads. This suite is used to demonstrate speed up scalability of SMPs and is cluster scalable. The most significant influence on scalability on this benchmark is the parallelizing compiler and the workload. OMP2001 contains two benchmark suites: OMPM2001 for use with mid-range server, those having less than 32 processors, and OMPL2001 which measures performance of high-end servers with up to 512 processors. A medium data set and large data set of the same 91 applications are used within the two benchmark suites along with an additional FORTRAN90 and an additional C code for the medium version.

SPECjbb2000 is a CPU intensive benchmark used to evaluate the performance of hardware and software aspects of Java Virtual Machine (JVM) servers. The workload represents an order processing application for a wholesale supplier written in Java. This benchmark and problem size were designed for small to mid-range SMPs and 32-bit JVMs but is now used with both

32-bit and 64-bit JVM implementation. **SPECjbb2005** is similar to JBB2000 with changes to the database internal structure, system garbage collector calls, and data handling.

LINPACK is a collection of Fortran subroutines that analyze and solve linear equations and linear least-squares problems normally used in an HPC environment. It was developed by Jack Dongarra, Jim Bunch, Cleve Moler and Pete Stewart. The package solves linear systems whose matrices are general, banded, symmetric indefinite, symmetric positive definite, triangular, and tri-diagonal square. In addition, the package computes the QR and singular value decompositions of rectangular matrices and applies them to least-squares problems. Matrix sizes can be "n x n". The "n x n" is user defined which might involve running multiple CPUs.

IBM Corporation
Marketing Communications
Systems Group
Route 100
Somers, New York 10589

Produced in the United States of America
October 2005

The IBM home page on the Internet can be found at http://www.**ibm.com**.

The System p5 and @server p5 home page on the Internet can be found at http://www.**ibm.com**/servers/eserver/pseries.