



IBM @server p5 510 and p5 510 Express offer IBM POWER5 technology in a 2U rack drawer package

Overview

IBM @server p5 510 and IBM @server p5 510 Express Rack-mount Servers (9110-510) give you new tools for managing On Demand Business, greater application flexibility, and innovative technology—all designed to help you capitalize on the On Demand Business revolution.

The p5-510 Express comes in a 2U rack drawer package. It is available in a 1-way or 2-way configuration using state-of-the-art, 64-bit, copper-based POWER5™ microprocessors running at 1.5 GHz.

The p5-510 is very similar to the p5-510 Express server except that the POWER5 microprocessor operates at 1.65 GHz. To simplify, both products are referred to as p5-510 except where a distinction needs to be made between the Express offering and the non-Express offering.

The new POWER5 microprocessor provides many performance and reliability advances over the POWER4™ architecture. Chief among the enhancements is Simultaneous Multi-Threading (SMT), which is designed to significantly boost performance by allowing two code streams to concurrently access the several execution units on the processor.

The p5-510 has a base of 512 MB of main memory that can be expanded to 32 GB for faster performance and exploitation of 64-bit addressing as used in large database applications.

The p5-510 contains five bays. Four of the five are front-accessible, hot-swap-capable DASD bays and can accommodate up to 587.2 GB of disk storage. The fifth bay is used for a slimline DVD-ROM or DVD-RAM.

Other integrated features include:

- Three 64-bit PCI-X slots

- Service Processor
- Two 10/100/1000 Ethernet ports, limited purpose serial ports, USB ports, and HMC ports
- Dual channel Ultra320 SCSI controller
- External SCSI port
- Hot-swap power and cooling
- Redundant power (optional) and cooling

Dynamic LPAR is supported on the p5-510, allowing up to two partitions. In addition, the optional Advanced POWER™ Virtualization feature supports up to 10 partitions per processor.

IBM Cluster Systems Management (CSM) V1.4 for AIX® 5L and Linux™ on POWER is supported on the p5-510. For hardware control, a Hardware Management Console (HMC) is required. Additionally, the p5-510 is added to the hardware models supported with the IBM @server pSeries® Cluster 1600 running CSM.

The p5-510 is backed by a three-year warranty.

Key Prerequisites

If installing AIX on the system (one of these):

- AIX 5L for POWER V5.2 with the 5200-04 Recommended Maintenance Package (APAR IY56722 plus APAR IY62266), or later
- AIX 5L for POWER V5.3 with the 5300-01 Recommended Maintenance Package (APAR IY62267), or later

If installing Linux on the system (one of these):

- SUSE LINUX Enterprise Server 9 for POWER, or later
- Red Hat Enterprise Linux AS for POWER Version 3, or later

At a Glance

IBM @server p5 510 and IBM @server p510 Express offer:

- Outstanding value for entry UNIX® servers
- Powerful, symmetric multiprocessing (SMP) on demand server
- 2U rack-mount configuration
- One 1-way or 2-way, state-of-the-art, 64-bit POWER5 processor
 - 1.5 GHz processor on the p5-510 Express
 - 1.65 GHz processor on the p5-510
- Performance and capacity needed by demanding on demand applications
- Up to 32 GB of memory
- Up to 587.2 GB of internal disk storage
- Four hot-swap disk bays and three PCI-X slots for feature expandability
- One slimline media bay for DVD-ROM or DVD-RAM
- Dynamic logical partitioning (LPAR)

Planned availability dates

April 22, 2005, for feature numbers 7612, 7449, and 8649.

February 18, 2005, for all other features.

This announcement is provided for your information only. For additional information, contact your IBM representative.

Description

IBM eServer p5 510 Value Paks

IBM eServer p5 510 and p5 510 Express

The standard features in the IBM eServer p5 510 and p5 510 Express include:

- Rack-mount (2U) configuration
- 1-way and 2-way SMP design (one processor card)
- POWER5 64-bit processor
 - p5-510 Express
 - 1-way 1.5 GHz with no L3 cache
 - 2-way 1.5 GHz with 36 MB L3 cache (available April 22, 2005)
 - p5-510
 - 1-way 1.65 GHz with 36 MB L3 cache
 - 2-way 1.65 GHz with 36 MB L3 cache
- 512 MB of DDR-I 266 MHz ECC memory
 - Expandable to 32 GB
- Four hot-swap disk drive bays
- One slimline media bay:
 - One DVD-ROM (optional — default)
 - One DVD-RAM (optional)
- Three PCI-X slots:
 - 64-bit, 3.3 volt, 133 MHz (long)
- Integrated:
 - Dual 10/100/1000 Ethernet
 - Dual Channel Ultra320 SCSI controller
 - One external SCSI port
 - Service Processor
 - Hot-swap and redundant fans
 - Two USB ports
 - Two limited purpose serial ports
 - Two HMC ports
- 700 watt power supply, base (redundant power optional)
- AIX 5L V5.2 or V5.3, or later, operating system, or SUSE LINUX Enterprise Server 9 for POWER, or later, or Red Hat Enterprise Linux AS for POWER Version 3, or later.

The minimum p5-510 configuration must include a processor, a processor entitlement, memory, power supply, HDD, a bezel and hardware indicator, a power cord, and a language group specify. The IBM/OEM Rack-mount Drawer Rail Kit (#7166) is not part of the minimum configuration. The defaults, if no choice is made, are:

1-way Processor, 1.5 GHz	(#7609)
One processor entitlement	(#7442)
512 MB memory DIMM	(#4443)
Power supply, 700 watt, base	(#7989)
73.4 GB hard disk drive	(#3274)
IBM bezel and hardware	(#7998)
Language group specify	(#9300 or 97xx)
Power cord	(#6458, #6460, #6469-#6479, #6487, #6488, #6493-#6496, #6660, #6669, #6670, #6680, #6687)

A keyboard, mouse, graphics adapter, and monitor are available as options.

Value Paks are only available as initial order.

If you order a p5-510 Express with a 1-way POWER5 1.5 GHz processor (#7609) or a 2-way POWER5 1.5 GHz processor (#7612), or a p5-510 with a 1-way POWER5 1.65 GHz processor (#7610), or a 2-way POWER5 1.65 GHz processor (#7611) and a minimum configuration of memory, disks, and media device, you will qualify for a processor entitlement at no additional charge. The number of processors, total memory, quantity/size of disk, and presence of a media device are the only features that determine if a customer is entitled to processor entitlements at no additional charge. You will be entitled to a discounted AIX operating system license or may choose to purchase the system with no operating system. The discounted AIX operating system license is processed via a feature number on AIX.

The minimum criteria for processor entitlements at no additional charge are 1 GB of memory per active processor, two 73.4 GB hard disk drives, and a DVD-ROM. A DVD-RAM may be substituted for the DVD-ROM.

You may make changes to the standard features as needed and still qualify for processor entitlements at no additional charge and a discounted AIX operating system license. Selection of memory or disk drives smaller than those defined as the Value Pak minimums will disqualify the order as a Value Pak.

Express Value Pak 1

1-way 1.5 GHz processor (1 x #7609)
1 GB memory (Default: 1 x #4444)
2 x 73.4 GB disk drive (#3274)
1 DVD-ROM (#2640)

Receive one processor entitlement (#8642) at no additional charge. One processor is entitled.

Express Value Pak 2 (available on April 22, 2005):

2-way 1.5 GHz processor (1 x #7612)
2 GB memory (Default: 1 x #4447)
2 x 73.4 GB disk drive (#3274)
1 DVD-ROM (#2640)
One processor entitlement (#7449)

Receive one processor entitlement (#8649) at no additional charge. Two processors are entitled.

Value Pak 3:

1-way 1.65 GHz processor, 36 MB L3 cache (1 x #7610)
1 GB memory (Default: 1 x #4444)
2 x 73.4 GB disk drive (#3274)
1 DVD-ROM (#2640)

Receive one processor entitlement (#8643) at no additional charge. One processor is entitled.

Value Pak 4:

2-way 1.65 GHz processor (1 x #7611)
2 GB memory (Default: 1 x #4447)
2 x 73.4 GB disk drive (#3274)
1 DVD-ROM (#2640)
One processor entitlement (#7444)

Receive one processor entitlement (#8644) at no additional charge. Two processors are entitled.

Reliability, availability, and serviceability (RAS)

Reliability, fault tolerance, and data integrity: The reliability of the model 510 system starts with components, devices, and subsystems that are designed to be fault-tolerant. During the design and development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing, systems go through a thorough testing process designed to help ensure the highest level of product quality.

The model 510 system L3 cache and system memory offers ECC (error checking and correcting) fault-tolerant features. ECC is designed to correct environmentally induced, single-bit, intermittent memory failures and single-bit hard failures. With ECC, the likelihood of memory failures will be substantially reduced. ECC also provides double-bit memory error detection that helps protect data integrity in the event of a double-bit memory failure. System memory also provides 4-bit packet error detection that helps to protect data integrity in the event of a DRAM chip failure. The system bus, I/O bus, and PCI buses are designed with parity error detection.

Disk mirroring and disk controller duplexing are also provided by the AIX operating system. Linux supports DASD mirroring (RAID 1). This is supported in software using the md driver. Some of the hardware RAID adapters supported under Linux also support mirroring.

The Journaled File System maintains file system consistency and reduces the likelihood of data loss when the system is abnormally halted due to a power failure.

PCI extended error handling (EEH)¹: In the past, PCI bus parity errors caused a global machine check interrupt, which eventually required a system reboot to continue. In the POWER5 systems, new I/O drawer hardware, system firmware, and AIX 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. This mechanism is called PCI EEH. EEH-enabled adapters respond to a special data packet generated from the affected PCI 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. Currently, Linux does not support the EEH behavior.

¹ Indicates that this RAS function is not supported under Linux.

Memory error correction extensions: The standard memory has single-error-correct and double-error-detect ECC circuitry designed to correct single-bit memory failures. The double-bit detection is designed to help maintain data integrity by detecting and reporting multiple errors beyond what the ECC circuitry can correct. The memory chips are organized such that the failure of any specific memory module only affects a single bit within an ECC word (bit scattering), thus allowing for error correction and continued operation in the presence of a complete chip failure (Chipkill™ recovery).

The memory also utilizes memory scrubbing and thresholding to determine when spare memory modules, within each bank of memory, if available, should be used to replace ones that have exceeded their threshold value (dynamic bit steering).

Redundancy for array self-healing: Although the most likely failure event in a processor is a soft single-bit error in one of its caches, there are other events that can occur, and they need to be distinguished from one another. For the L1, L2, and L3 caches and their directories, hardware

and firmware keep track of whether permanent errors are being corrected beyond a threshold. If exceeded, a deferred repair error log is created. Additional run-time availability actions, such as CPU vary off² or L3 cache line delete, are also initiated.

L1 and L2 caches and L2 and L3 directories on the POWER5 chip are manufactured with spare bits in their arrays that can be accessed via programmable steering logic to replace faulty bits in the respective arrays. This is analogous to the redundant bit steering employed in main storage as a mechanism that is designed to help avoid physical repair, and is also implemented in POWER5 systems. The steering logic is activated during processor initialization and is initiated by the built-in self-test (BIST) at power-on time.

L3 cache redundancy is implemented at the cache line level. Exceeding correctable error thresholds while running causes a dynamic L3 cache line delete function to be invoked.

² Indicates that this RAS function is only available for a Linux operating system running the 2.6 kernel.

Service Processor: The Service Processor included in the model 510 gives you an immediate means to diagnose, check status, and sense operational conditions of a remote system, even when the main processor is inoperable.

The Service Processor enables firmware and operating system surveillance, several remote power controls, environmental monitoring (only critical errors are supported under Linux), reset, boot features, remote maintenance, and diagnostic activities, including console mirroring. The Service Processor can place calls to report surveillance failures, critical environmental faults, and critical processing faults.

Fault monitoring functions

- BIST (built-in self-test) and POST (power-on self-test) check the processor, L3 cache, memory, and associated hardware required for proper booting of the operating system every time the system is powered on. If a noncritical error is detected or if the errors occur in the resources that can be removed from the system configuration, the booting process is designed to proceed to completion. The errors are logged in the system nonvolatile RAM (NVRAM).
- Disk drive fault tracking can alert the system administrator of an impending disk failure before it impacts customer operation.
- The AIX or Linux log (where hardware and software failures are recorded and analyzed by the Error Log Analysis [ELA] routine) warns the system administrator about the causes of system problems. This also enables IBM service representatives to bring along probable replacement hardware components when a service call is placed, thus minimizing system repair time.

Mutual surveillance: The Service Processor monitors the operation of the POWER Hypervisor firmware during the boot process and watches for loss of control during system operation. It also allows the POWER Hypervisor to monitor Service Processor activity. The Service Processor can take appropriate action, including calling for service, when it detects the POWER Hypervisor firmware has lost control. Likewise, the POWER Hypervisor can request a Service Processor repair action if necessary.

First Failure Data Capture, a full-spectrum diagnostics strategy: Diagnosing problems in a computer is a critical requirement for autonomic computing. The first step to producing a computer that truly has the ability to self-heal is to create a highly accurate way to identify and isolate hardware errors. IBM has implemented a server design that builds in hardware error-check stations that capture and help to identify error conditions within the server. Each of these checkers is viewed as a diagnostic probe into the server, and, when coupled with extensive diagnostic firmware routines, allows quick and accurate assessment of hardware error conditions at run-time.

FFDC check stations are carefully positioned within the server logic and data paths to help ensure that potential errors can be quickly identified and accurately tracked to an individual field replaceable unit (FRU). These checkers are collected in a series of Fault Isolation Registers, where they can easily be accessed by the Service Processor. All communication between the SP and the FIR is accomplished "out of band." That is, operation of the error-detection mechanism is transparent to an operating system. This entire structure is "below the architecture" and is not seen, nor accessed, by system-level activities.

Environmental monitoring functions

- Temperature monitoring increases the fan speed rotation when ambient temperature is above the normal operating range.
- Temperature monitoring warns the system administrator of potential environmental related-problems (for example, air conditioning and air circulation around the system) so that appropriate corrective actions can be taken before a critical failure threshold is reached. It also performs an orderly system shutdown when the operating temperature exceeds the critical level.
- Fan speed monitoring provides a warning and an orderly system shutdown when the speed is out of the operational specification.
- Voltage monitoring provides a warning and an orderly system shutdown when the voltages are out of the operational specification.

Error handling and reporting: In the unlikely event of system hardware or environmentally induced failure, the system run-time error capture capability systematically analyzes the hardware error signature to determine the cause of failure. The analysis will be stored in the system NVRAM. When the system can be successfully rebooted either manually or automatically, the error will be reported to the AIX or Linux operating system. ELA can be used to display the failure cause and the physical location of failing hardware.

With the integrated Service Processor, the system has the ability to automatically send out an alert via phone line to a pager or call for service in the event of critical system failure. A hardware fault will also turn on the two Attention Indicators (one located on the front of the system unit and the other on the rear of the system) to alert the user of an internal hardware problem. The indicator may also be turned on by the operator as a tool to allow system identification. For identification, the indicators will flash, whereas the indicator will be on solid when an error condition occurs.

Availability enhancement functions

- The auto-restart (reboot) option, when enabled, can reboot the system automatically following an unrecoverable software error, software hang, hardware failure, or environmentally induced (ac power) failure.

Serviceability: Model 510 is designed for customer setup of the machine and for the subsequent addition of most features (adapters/devices)². For a fee, IBM Service can perform the installation.

- Model 510 allows customers to replace service parts (customer replaceable unit) if they want to. Model 510 has incorporated LEDs that will indicate the parts needing to be replaced.
- Model 510 allows support personnel to remotely log into a system to review error logs and perform remote maintenance. The model 510 Service Processor enables the analysis of a system that will not boot.
- The diagnostics consist of Stand-alone Diagnostics, which are loaded from the DVD-ROM drive, and Online Diagnostics.
- Online Diagnostics¹, when installed, are resident with AIX on the disk or system. 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 Error Log and the AIX Configuration Data.
 - Service mode allows checking of system devices and features.
 - Maintenance mode allows checking of most system resources.
 - Concurrent mode allows the normal system functions to continue while selected resources are being checked.
- 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 9110-510 system has an optional DVD-ROM (#2640) and DVD-RAM (#5751), alternate methods for maintaining and servicing the system need to be available if the DVD-ROM or DVD-RAM is not ordered; an external Internet connection must be available to maintain or update system microcode to the latest required level.

Service Agent: The Service Agent is available at no additional charge. When installed on an @server system, the Service Agent can enhance IBM's ability to provide the system with maintenance service.

The Service Agent:

- Monitors and analyzes system errors and, if needed, can automatically place a service call to IBM without customer intervention
- Can help reduce the effect of business disruptions due to unplanned system outages and failures
- Performs problem analysis on a subset of hardware-related problems and, with customer authorization, can report automatically the results to IBM Service

Online customer support (OCS)¹: OCS for hardware problem reporting may be performed via remote login by eServer specialists. The Electronic Service Agent™ software can also be used for this capability.

AIX Support offerings will be under AIXSERV and Electronic Service Agent.

Statement of general direction

IBM intends to provide High Availability Cluster Multi-Processing (HACMP™) V5.2 support on the p5-510 running AIX 5L V5.2 and V5.3 by second quarter 2005.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

Trademarks

POWER5, POWER4, POWER, Chipkill, Electronic Service Agent, HACMP, EIA, and ThinkVision are trademarks of International Business Machines Corporation in the United States or other countries or both.

eServer, the e-business logo, AIX, pSeries, TotalStorage, and RS/6000 are registered trademarks of International Business Machines Corporation in the United States or other countries or both.

UNIX is a registered trademark of the Open Company 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, and service names may be trademarks or service marks of others.