IBM @server® p5 575 cluster node

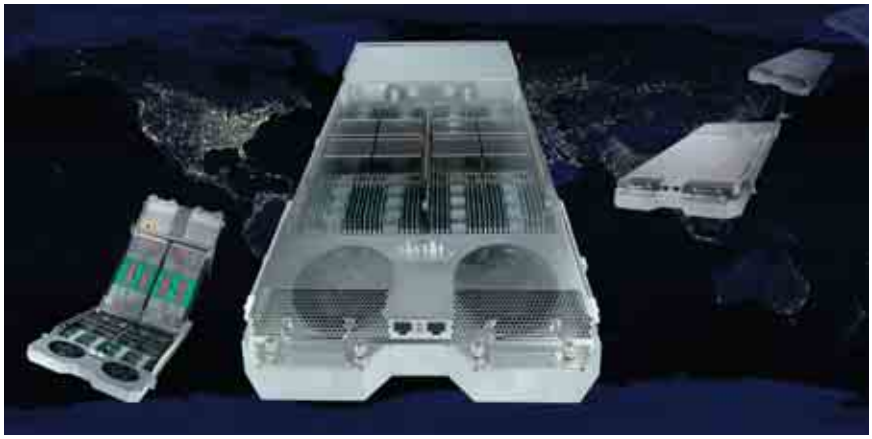

@server p5 575 cluster node

Highlights

- ***High-powered building block designed for the most demanding HPC and BI applications***
- ***Innovative packaging helps to minimize space requirements, reduce environmental demands, and simplify system management***
- ***Can be configured with up to 128 nodes in a single Cluster 1600 system—a 1024 CPU supercomputer***

The adage "time is money" defines today's supercharged, high-performance economy. Companies involved in drug design, oil reservoir modeling, weather forecasting, financial simulation and business intelligence (BI) know that wasted time can easily result in missed opportunities and the loss of competitive advantage. As a result, businesses, research laboratories and academic institutions need state-of-the-art, high-performing systems that can speed discovery, reduce time to market and farm enterprise data faster—all while keeping costs under control.

The IBM @server® p5 575 cluster node is specifically designed to tackle these "extreme" performance computing applications, which require both high computational performance and memory bandwidth. The p5-575 node includes eight powerful 1.9 GHz IBM POWER5 microprocessors. Each processor includes 1.9MB of L2 and 36MB of L3 dedicated cache for the

ultimate in high-bandwidth computing. No POWER5 processor-based system can match the extraordinary density achieved by 12 p5-575 cluster nodes packaged in a single 24-inch system frame. Compared to its predecessor, the IBM @server pSeries® 655, the p5-575 delivers substantially higher performance for memory bandwidth-intensive applications.

p5-575: does more than crunch numbers

While other cluster nodes are geared solely toward fast computations, the p5-575 does much more than number crunching. It is designed to meet the needs of a broad array of organizations that require not only fast processing but also rapid and continuous access to vast amounts of data. With nearly 100 GBps of peak memory bandwidth per node, the p5-575 can help organizations conduct oceanographic studies, weather observations, computational fluid dynamics, energy research and other bandwidth-intensive work that requires transferring, accessing and rapidly analyzing large quantities of data.

Like the p655, the p5-575 is designed to be an excellent match for many businesses, such as insurance, banking, finance and retail organizations that have amassed large quantities of information and want to mine that data for competitive advantage. Many of the features that make the p5-575 an excellent fit for the most demanding engineering and scientific tasks also makes it well-suited for large-scale data warehousing and data servicer applications using IBM DB2® Universal Database™ software for BI. The node can be configured so that it only need be replicated to scale-out a cluster.

Advanced processor technology delivers outstanding performance

The p5-575 is designed to deliver exceptional performance with its eight 64-bit POWER5 processors. Those processors incorporate simultaneous multi-threading[1], which allows two application threads to be executed concurrently. The result is improved performance compared with earlier IBM POWER™ processor–based systems.

Data I/O handling performance is greatly assisted by the L2 and L3 caches available to each processor. These caches help to stage information more effectively from processor memory to applications, allowing the p5-575 cluster node to run workloads significantly faster than its predecessor.

To further enhance system performance, the p5-575 memory DIMMs have eight point-to-point connections to each of the eight processors, with a maximum memory capacity of 32GB per processor and a peak memory transfer speed of up to 12.4 GBps per processor—or nearly 100 GBps for an 8-way node. The DIMMs are in close proximity to supported processor cores, reducing signal propagation delay and lowering power and heat dissipation requirements.

When compared with smaller symmetric multiprocessing (SMP) cluster nodes deployed in high performance computing (HPC) clusters, the p5-575 enables a greater proportion of the workload to communicate over a lightning-fast, low-latency, high-bandwidth SMP fabric, as

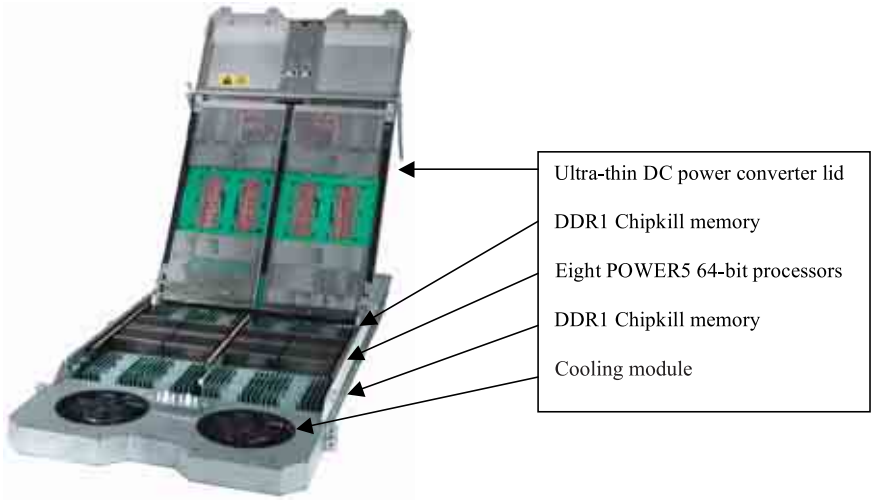

opposed to an I/O-based switch fabric. This system configuration can deliver significantly better overall system performance while reducing complexity, improving manageability and helping to contain costs.

As an option, the p5-575 offers advanced IBM Virtualization Engine™ system technologies[1] with Micro-Partitioning™ capabilities to help optimize the use of system resources. With the IBM Virtualization Engine, multiple copies of operating systems can be run on the same processor, helping reduce the number of cluster nodes and minimize software licensing costs. Micro-Partitioning technology allows processors to be finely divided (up to ten micro-partitions per processor) so that more work can be executed on a single processor. At the same time, IBM logical partitioning (LPAR) technology ensures that application data running on one partition is shielded from data on another partition, providing a high level of data security, integrity and increased application availability.

Innovative design minimizes floor space and enhances reliability

The p5-575 cluster node features innovative, elegant design and packaging. Mounted in a sleek 2U enclosure, the modular p5-575 allows users to deploy up to 12 nodes in a single 42U system frame.

Nodes can be configured with or without support for optional internal and external I/O devices. The streamlined "computational" node includes two dual 10/100/1000 Mbps Ethernet ports; two integrated Ultra3 SCSI controllers; two Hardware Management Console (HMC) ports for system control, logical partitioning and service functions; and two hot-swappable disk storage bays, which accommodate 10K rpm or 15K rpm disk drives. The "I/O" node adds four 133 MHz hot-plug, blind-swap PCI-X adapter slots and an RIO-2 hub port to attach an optional I/O drawer.

In both configurations, the unique node enclosure has four component modules, each one custom-designed to satisfy the requirements of high-performance, high-density computing. Supplementing the I/O module, the highly efficient DC power distribution module is integrated into the lid of the node. This innovative power system relies on embedded circuitry rather than external wiring, providing more reliable and efficient power distribution. The hinged lid opens easily for access to the processor and memory module, which contains the POWER5 processors and system memory DIMMs.

The processor and memory module is the heart of the system. The p5-575 features eight POWER5 processors, each with 36MB of dedicated L3 cache and point-to-point connections to up to eight memory DIMMs. This implementation helps to provide exceptionally high memory bandwidth to support many demanding HPC applications.

The front-end cooling module has two air-intake ventilation grids and two custom-designed blowers with high-capacity impellers and high-efficiency motors that are designed for extended life and easy serviceability.

Scale-up or out easily and inexpensively

The p5-575 cluster node can be scaled within the enclosure or replicated within the cluster to meet growing workload requirements. Equipped with 1GB of DDR1 memory in its minimum configuration, each node can scale-up to 256GB. Two hot-swappable disk drives allow disk storage capacity from 72.8GB to 293.6GB. For even greater disk capacity, the I/O node configuration supports a 4U I/O drawer through a RIO-2 hub port at the back of the enclosure. Each I/O drawer holds up to 16 additional disk storage bays, accommodating up to 1.17TB of disk storage. Two cluster nodes can share a single I/O drawer, with each system frame containing up to five I/O drawers.

The p5-575 can scale-out easily and cost-effectively as workload requirements increase. Each system frame accommodates up to 12 p5-575 cluster nodes. Organizations can add system frames to build a system cluster with anywhere from 16 to more than 512 processors. Cluster nodes are attached with industry standard Ethernet (10/100 Mbps or 1 Gbps). Support of the IBM @server pSeries High Performance Switch is planned for the second quarter of 2005.[2]

Mainframe-inspired RAS features help to provide peace of mind

Although the p5-575 cluster node comes in a small package, it is loaded with mainframe-inspired features that help to ensure high reliability, availability and serviceability (RAS). The p5-575 is equipped with a built-in service processor which monitors system operations continuously and can take preventive or corrective action for quick problem resolution. First Failure Data Capture (FFDC) capabilities help to identify and log problems before system failures occur. IBM error checking and correction (ECC) / Chipkill™ memory technology detects and corrects memory errors to help prevent costly system crashes. Finally, Dynamic Processor Deallocation capabilities in many cases can identify potential processor problems, generate error reports and deallocate processors before they fail.

The p5-575 system includes structural elements to help ensure outstanding availability. Redundant bulk power and optional internal battery back-up help administrators keep the nodes running even in the event of power problems. In the event a cooling fan fails, the second

fan will increase its velocity and the system service processor may initiate a service call.

The p5-575 power distribution and conversion system—adopted from the @server p5 595 server design—relies on embedded circuitry rather than external wiring to distribute power among system components with the objective of providing more reliable and efficient power distribution. In addition, the p5-575 uses IBM's leading-edge rack level distributed power conversion architecture to maximize system density, simplify power connection and provide a robust, redundant system power supply arrangement. Two simple, neutral free universal line cords connect the p5-575 system frame to a client's facility anywhere in the world with no adjustments being required to personalize for power utility voltage or frequency. Support for 200v to 240v, 380v to 415v, and 480v three phase power inputs should allow clients to enjoy reduced facility equipment cost and help improve energy efficiency. The ability of the p5-575 to tolerate power disturbances is exceptional in comparison to most other computing equipment, and optional battery back-up can help the system ride through a momentary power interruption without the need for large and expensive Universal Power Supply (UPS) systems.

Built-in reliability features

IBM autonomic computing enhancements are built into the p5-575. Self-protecting helps the p5-575 determine the cause of an error as it happens and may reduce lengthy service times attempting to recreate errors after the fact. Errors may be self-correcting or resources varied off-line while the server remains available for use. IBM's First Failure Data Capture provides error information in real-time and makes it possible to determine the parts needed to fix the problem. The service processor has the capability to determine which part or component needs repair and initiate a service call to identify parts needed for maintenance at a time acceptable to the client.

Self-healing capabilities help the p5-575 to overcome error conditions and continue operating if a failure is detected. This is implemented through Error Checking and Correcting Code (ECC) L2 and L3 caches and main memory and through bit-scattering, bit-steering and memory scrubbing soft-error recovery procedures in main memory. Bit-scattering scatters bits across four different memory words, enables recovery of single-bit errors and should keep a p5-575 running when a failure is detected by Chipkill memory. Bit-steering dynamically routes a bit to a spare memory chip in the event the memory failure rate for the bit exceeds a given threshold. If all bits should become used up on the spare chip, the service processor is invoked to request deferred maintenance at a time acceptable to the client. Memory scrubbing for soft single-bit errors is performed in the background so as to correct errors while memory is idle. This helps to prevent multiple-bit errors.

The p5-575 supports business-critical applications

The p5-575 cluster node can run AIX 5L™ and Linux® operating systems (OS) on the same node simultaneously, to provide the flexibility to support a full range of applications.

AIX 5L is IBM's industrial-strength UNIX® environment specially tuned for application performance and loaded with exceptional RAS features. The AIX 5L OS delivers enhancements to Java™ technology, Web performance and scalability for managing clusters of all sizes. Web-based remote management tools give administrators centralized control of the system, enabling them to monitor key resources, including adapter and network availability, file system status and processor workload.

The AIX 5L OS also incorporates Workload Manager, a resource management tool that specifies the relative importance of workloads to balance the demands of competing workloads and enhance system resources. Workload Manager can help ensure that applications remain responsive even during periods of peak system demand.

By supporting the Linux OS, the p5-575 cluster node offers important cost-saving opportunities. Because Linux is an open source technology, it is less expensive to license than many proprietary operating systems.

Nevertheless, Linux does not compromise on functionality. With a growing list of Linux applications available, it provides the freedom to use the right applications for organizations' needs. The Linux OS is available from IBM and selected Linux distributors in packages that include a range of open source tools and applications. IBM is firmly committed to Linux and offers expert service and support.

Software tools facilitate easy cluster management

IBM **@server** Cluster 1600 is a highly scalable cluster solution for UNIX or Linux environments, including the AIX 5L V5.2, AIX 5L V5.3, SUSE Linux Enterprise Server 9 (SLES 9) and Red Hat Enterprise Linux AS 3 (RHEL AS 3) operating systems. Cluster 1600 is implemented through Cluster Systems Management (CSM) for AIX 5L or Linux clusters. CSM supports other optional cluster software for HPC including:

- *Parallel Environment (PE) for AIX 5L–a high function development and execution environment for parallel message-passing applications.*
- *LoadLeveler®–dynamic job scheduling and workload balancing software supporting thousands of jobs within the cluster. LoadLeveler is supported on AIX 5L V5.2, V5.3 and SLES 9.*
- *GPFS–a high-performance, shared disk file system providing fast data access to all nodes in a cluster. GPFS is supported on AIX 5L V5.2, V5.3 and SLES 9.*
- *ESSL and Parallel ESSL–mathematical libraries for both AIX 5L and Linux to enhance performance of serial, parallel and scientific applications. Parallel ESSL is supported on AIX 5L V5.2, V5.3 and SLES 9, while ESSL is supported on AIX 5L V5.2, V5.3, SLES 9 and RHEL AS 3.*
- *High Availability Cluster Multiprocessing (HACMP™) for AIX 5L–helps provide continuous access to data and applications through database or application failover to a secondary server if the database or application server fails.*

Major productivity enhancements are provided through the POWER Hypervisor™ in conjunction with newly available operating systems. The user can establish dynamic logical partitions (dynamic LPAR) running AIX 5L V5.2, AIX 5L V5.3 or SLES 9 operating systems. Dynamic LPAR enables system administrators to reallocate system resources without rebooting the system or the partition.

If AIX 5L V5.3, SLES 9 or RHEL AS 3 are selected for a partition, the user can take advantage of the benefits of hardware simultaneous multi-threading, which may provide an increase of up to 30% (based on rPerf projections[3]) in processor throughput over single-threaded operation, depending on the nature of the applications being run on the partition. Furthermore, with AIX 5L V5.3 or SLES 9, the user can obtain even more flexibility with the optionally available Advanced POWER Virtualization, which provides Micro-Partitioning™, shared processor pool and Virtual I/O capabilities.

Micro-Partitioning provides the capability to establish up to 80 LPARs on a single p5-575 node, effectively splitting each processor's power among up to 10 LPARs. Shared processor pool provides a pool of processing power that is shared among partitions to improve utilization and throughput, and which can be changed dynamically to meet changing environments. Virtual I/O enables the physical sharing of disk drives and communications adapters and helps reduce the number of expensive devices and improve system administration and utilization. It also enables high-speed, secure partition-to-partition communication to help improve performance.

An additional capability of Advanced POWER Virtualization supported by AIX 5L V5.3 is Partition Load Manager which provides policy-based, automatic partition resource tuning that can adjust CPU and memory allocations in response to load demands.

The p5-575 helps to build an infrastructure for the future

The IBM @server p5 575 cluster node is designed to be a high-performance building block for supercomputing. With high sustained throughput and I/O bandwidth, it is an excellent match for scientific and engineering HPC applications as well as BI functions for which organizations need to transfer, access and analyze large amounts of data rapidly. The innovative, highly dense packaging of the p5-575 cluster node, its optional IBM Virtualization Engine with Micro-Partitioning capabilities and its ability to run the AIX 5L and Linux operating systems concurrently, will help get more work done while using less physical floor space. The comprehensive set of cluster management tools designed for the AIX 5L and Linux operating systems help provide the means to assemble and effectively manage a cluster. And with easy scalability, the p5-575 will be ready to grow with an organization's high-performance requirements.

p5-575 At a glance

Minimum configuration	
Microprocessors	8-way SMP 1.9 GHz POWER5 microprocessor
L2 cache	15.2MB/cluster node
L3 cache	288MB/cluster node
RAM (memory)	1GB of 266 MHz DDR1
Disk bays	Two hot-swappable Ultra320 disk bays, up to 293.2GB of storage capacity (73.4/146.8GB 10K rpm or 36.4/73.4GB 15K rpm disk drives available)
Expansion slots	Four 133 MHz PCI-X
PCI-X adapter support	32- and 64-bit
Standard features	
I/O ports	Two Ultra3 SCSI controllers I/O node configuration: • Four 10/100/1000 Mbps Ethernet ports • Two HMC ports • RIO-2 hub port for optional I/O drawer Compute node configuration: • Four 10/100/1000 Mbps Ethernet ports • Two HMC ports
System expansion	
RAM	Up to 256GB DDR1
I/O drawer	One—can be shared by two cluster nodes
Expansion slots	20 PCI-X via optional I/O drawer
Disk bays	Up to 16 hot-swappable disk bays via optional I/O drawer; up to 1.17TB of additional storage capacity (36.4GB/73.4GB 15K rpm disk drives available)
Logical partitioning support	Dynamic LPAR*
IBM Virtualization Engine (optional)	Micro-Partitioning Shared processor pool Virtual LAN Virtual I/O
Battery backup	Up to six—redundant or non-redundant

* Available with AIX 5L and SLES 9 operating system

** With slim-line doors and populated with 12 p5-575, internal battery backup and one I/O drawer. Weight will vary when disks, adapters and other peripherals are installed.

p5-575 At a glance

RAS features	Copper, silicon-on-insulator (SOI) microprocessors IBM Chipkill ECC, bit-steering memory ECC L2 cache, L3 cache Hot-swappable disk bays Hot-plug/blind-swap PCI-X slots Hot-plug power supplies and cooling fans Redundant bulk power supply Dynamic Processor Deallocation Dynamic deallocation of logical partitions and PCI-X bus slots Battery backup and redundant battery backup (optional)
Operating systems	AIX 5L Versions 5.2/5.3 Linux • SUSE LINUX Enterprise Server 9 for POWER (SLES 9) or later • Red Hat Enterprise Linux AS 3 for POWER Update 4 (RHEL AS 3) or later
Power requirements	200v to 240v; 380v to 415v; 480v AC
System frame dimensions	79.7"H x 30.9"W x 60.2"D (202.5 cm x 78.5 cm x 153.0 cm); weight: 3,460 lb (1,573 kg) **
Warranty	On site, 8 A.M.—5 P.M. next-business-day for one year

* Available with AIX 5L and SLES 9 operating system

** With slim-line doors and populated with 12 p5-575, internal battery backup and one I/O drawer. Weight will vary when disks, adapters and other peripherals are installed.

For more information

To learn more about the IBM @server p5 575 cluster node, contact your IBM marketing representative or IBM Business Partner, or visit the following Web sites:

- **ibm.com**/eserver/pseries
- **ibm.com**/servers/aix
- **ibm.com**/linux/power
- **ibm.com**/common/ssi

rPerf estimates are calculated based on systems with the latest levels of AIX 5L and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM @server pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration.

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. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

Many of the pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: **ibm.com**/servers/eserver/pseries/linux/whitepapers/linux_pseries.html.

[1] Not supported on AIX 5L V5.2.

[2] All statements regarding IBM's plans, directions and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

[3] rPerf (Relative Performance) is an estimate of commercial processing performance relative to other pSeries systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

IBM Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States
February 2005

PSD00762-USEN-01