IBM Power 575 supercomputing node

Power 575 node

Highlights

- ***Optimized for running large, highly parallel computationally-intensive workloads and algorithms***
- ***Application affinity for weather and climate modeling, computational chemistry, physics, computer-aided engineering, computational fluid dynamics and petroleum exploration***

The IBM Power™ 575 supercomputing node is designed for organizations that require a highly scalable system with extreme parallel processing performance and dense, modular packaging. Use it in clustered configurations of as few as 32 processor cores or in world-class supercomputer configurations of thousands of processors. Combined with specialized software from IBM, this system is designed to perform and represents the latest Power technology available.

Ideal workloads for this system include high performance computing (HPC) applications such as weather and climate modeling, computational chemistry, physics, computer-aided engineering, computational fluid dynamics and petroleum exploration that require highly intense computations where the workload is aligned with parallel processing methodologies. IBM has long been a leader in these application areas and, with this system, enables you to innovate and mold the future.

Densely packing up to 448 POWER6™ processor cores per frame, each one running at 4.7 GHz with innovative cooling features, the 32-core

Power 575 supercomputing node is designed for speed and tuned for performance. Available in increments of 2U building blocks, hundreds of these nodes can be clustered together to tackle the world's greatest problems. Supported by up to 3.5 TB of memory per frame and super fast interconnects, the Power 575 is estimated to achieve over five times the GFLOPS per frame as its predecessor with POWER5+™ technology.[1] The radical new approach to HPC represented by the Power 575 system marks another step in the evolution of modular clusters designed to solve the world's greatest problems. Starting with the lightning fast 4.7 GHz processor, adding the ultra dense packaging of 32 cores in a 2U node and injected with a chilled coolant to enable peak performance, this system is a supersonic race car on the IT highway.

Feature	Benefits
Extreme POWER6 performance	• Ability to solve problems faster or tackle even larger problems • Ability to respond to increasingly complex challenges • Ability to inject more variables into equations to improve modeling accuracy
Innovative modular design	• Ability to buy the exact size system required • Ability to easily grow the system over time • Greater packaging density requires less floor space thereby reducing operational costs
Radical energy-efficient cooling technology	• Enables greater performance with a smaller footprint • Supports environmental goals by optimizing energy usage • Reduces operational costs
Ultra-scalable	• Allows for extreme performance • Enables centralized management of multiple interconnected systems
Mainframe-inspired reliability	• Improved assurance that work will not be interrupted • Better customer satisfaction due to reliable and predictable operations • Faster repair when required due to sophisticated system diagnostics
Extensive software support	• Ability to select the right application to meet your needs

Extreme POWER6 performance

With years of patent leadership and decades of IT experience, it is no surprise that IBM technology is the industry-leader for processor performance. The ultra-high frequency of POWER6 575 technology provides a superscalar design with over twice the clock speed and twice the memory bandwidth of previous IBM generations. POWER6 processors feature enhanced simultaneous multithreading with support for Hardware Decimal Floating-Point and an AltiVec™ SIMD accelerator for added performance.

Arrange 32 of these high-powered engines together in a single slim-line node with 256 GB of memory and a dual 2-port 4X host channel adapter, and you have a system designed for speed. One Power 575 node is estimated to achieve over 600 GFLOPS, 4.95 times the performance of its predecessor with POWER5+ technology.[3]

Innovative modular design

Mounted in a sleek 2U enclosure, the Power 575 supercomputing node features innovative, elegant design and packaging that facilitates ease of service and flexibility. Few systems can match the extraordinary density achieved with 448 processor cores in a single footprint (fourteen 32-core nodes packaged in a single 42U 24-inch system frame). Nodes can be configured with or without support for internal and/or external I/O devices. Compared to its POWER5+ predecessor, the Power 575 node delivers twice the packaging density and thereby enables substantially higher sustained throughput for high performance computing (HPC) applications.

The highly efficient DC power distribution module is integrated into the hinged lid of the node. This innovative power system relies on embedded circuitry rather than external wiring, providing more reliable and efficient power distribution. This power module includes precision intelligent monitoring and control functions that are designed to help assure power delivery is optimized at all times and provides alert data to the node service processor in the case of a fault.

Radical energy-efficient cooling technology

With the introduction of the Power 575, IBM is announcing the availability of new water cooling technology that can help reduce related energy costs. Cold-plates are directly mounted over all 32 POWER6 processor cores and with the aid of pumping devices mounted in the base of the frame, direct chilled coolant over the processors to keep them running at peak intensity.

In addition, ambient air is directed across the remaining components and the entire air cooling load is reduced by the integrated rear door heat exchanger located at the rear of the system. Inside, sealed tubes filled with chilled water absorb the heat generated in a fully populated rack and carry it away so it is not released into the datacenter. Similar to the optional IBM Rear Door Heat eXchanger for rack installations, this unit is now integrated into the same cooling loop as the rest of the system and standard for all Power 575 systems.

Years in development by IBM, the industry's leader in IT patents, this revolutionary dual partnership of cold-plates and Rear Door Heat eXchanger are integrated into the advanced, innovative design of the Power 575 system. Leveraging existing chilled water supply for air conditioning systems already located in a majority of datacenters around the globe, this system is

designed to reduce server heat emissions by up to 80 percent (versus a comparable air cooled system providing equivalent system performance).[4]

Scale-up or scale-out easily

Power 575 supercomputing nodes can be easily scaled within the system frame or replicated within the cluster to meet growing workload requirements. Equipped with 32 GB of memory in its minimum configuration, each node can scale-up to 256 GB. Two hot-swappable disk drives allow disk storage capacity from 73.4 GB to 293.6 GB per node. For even greater disk capacity, the optional 12X I/O drawer supports up to 16 Ultra3 SCSI disk bays or the optional 12S expansion drawer for up to 12 hot-swap SAS drives per drawer.

A Power 575 cluster can scale-out easily and cost-effectively as workload requirements increase. Each system frame accommodates up to 14 nodes. Organizations can add system frames to build a system cluster with hundreds of nodes and thousands of processor cores. Each fully populated frame with 14 nodes and 448 processor cores is estimated to deliver over 8 TFLOPS in a single footprint.

Mainframe-inspired reliability features

Among the world-class RAS capabilities provided in the Power 575 are a sophisticated service processor; hot-plug, hot-swappable and redundant components; IBM Chipkill™ ECC and bit-steering memory; First Failure Data Capture mechanisms; and dynamic deallocation of system components. These capabilities help to increase system availability and allow more work to be processed with less operational disruption. New with POWER6 technology, Processor Instruction Retry provides for the continuous monitoring of processor status with the capability to restart a processor if certain errors are detected. If required, workloads can be redirected to alternate processors, all without disruption to application execution.

The Power 575 system includes structural elements at the frame level to help ensure outstanding availability even in the event of facility power problems. The 575 system frame uses IBM's leading-edge rack level distributed power conversion architecture to increase system density, simplify power connection and provide a robust, redundant system power supply arrangement. Support for 200v to 240v, 380v to 415v, and 480v three phase power inputs allow clients to enjoy reduced facility equipment cost and help improve energy efficiency. The ability of the 575 to tolerate power disturbances is exceptional in comparison to most other computing equipment.

Extensive software support

The Power 575 supercomputing node can run both AIX® and Linux® operating systems simultaneously to provide the flexibility to support a full range of HPC applications as required. AIX is an industrial-strength UNIX® operating system (OS) tuned for application performance and delivers exceptional scalability for 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 575 node also supports the Linux OS allowing a broad choice of applications to best fit client needs. Because Linux is an open source technology, it has the worldwide Linux community enhancing, contributing and validating

the Linux kernel. In choosing Linux, users can still take advantage of the many reliability and scalability features designed into the 575. And with a large list of open source, IBM and third party applications available, Linux offers the freedom to use the right applications and tools for an organization's needs.

Additional cluster management capabilities are available through the following optional IBM Power Systems Software products:

- *IBM Parallel Environment (PE)—a high function development and execution environment for parallel message-passing applications under AIX and Linux.*
- *Tivoli® Workload Scheduler LoadLeveler®—a job scheduler designed to maximize resource utilization and throughput in a cluster environment to get the most out of the resources.*
- *IBM General Parallel File System™ (GPFS™)—a high-performance, shared disk file system providing fast data access to all nodes in a cluster.*
- *IBM Engineering and Scientific Subroutine Library (ESSL) and IBM Parallel ESSL—mathematical libraries for both AIX and Linux to enhance performance of serial, parallel and scientific applications.*

PowerVM™ is the family of technologies, capabilities and offerings that deliver industry-leading virtualization on IBM POWER processor-based systems. On the Power 575, PowerVM includes base components provided with IBM Power Systems firmware which includes logical partitioning (LPAR) technologies. In addition, optional components and PowerVM Editions are designed to provide advanced virtualization technologies resulting in efficiencies in resource utilization and cost savings. These are managed through use of a hardware management console (HMC).

PowerVM Standard Edition includes Micro-Partitioning™ and Virtual I/O Server (VIOS) capabilities, which are designed to allow businesses to increase system utilization while helping to ensure applications continue to get the resources they need. VIOS allows for the sharing of disk and optical devices as well as communications and Fibre Channel adapters to help drive down complexity and systems/administrative expenses. Also included is support for Multiple Shared Processor Pools, which allows for automatic non-disruptive balancing of processing power between partitions assigned to the shared pools, and Shared Dedicated Capacity, which helps optimize use of processor cycles.

PowerVM Enterprise Edition includes all the features of Standard Edition plus Live Partition Mobility (LPM), which allows a partition to be relocated from one server to another with virtually no impact to the applications running inside the partition. LPM is designed to enable servers to work together to help optimize system utilization and energy savings, improve application availability, balance critical workloads across multiple systems and respond to ever-changing business demands.

Power 575 at a glance

Processor cores	32 4.7 GHz POWER6 processor cores per node
Cache	4 MB L2 cache per processor core 32 MB L3 cache shared per two cores
RAM (memory)	Up to 256 GB per node
Internal disk	Two SAS small form factor disks per node (73.4 GB or 146.8 GB 10K rpm)
I/O	One integrated SAS controller for internal disks One integrated SAS connector for SAS disk drawer connection Two integrated dual-ported 10/100/1000 Ethernet ports Optional dual-port 10 Gigabit Optical Ethernet Four optional PCIe adaptor slots (two PCI risers each with two PCIe adapters) Optional dual 2-port 4x Host Channel Adapter (takes one PCI slot) Optional I/O drawer providing 20 blind-swap 64-bit PCI-X slots and up to 16 Ultra3 SCSI disk bays Optional I/O drawer providing 12 hot-swap SAS drives
PowerVM virtualization technologies	
POWER Hypervisor™	Dynamic LPAR; Virtual LAN (Memory to memory inter-partition communication)
PowerVM Standard Edition (optional)	Micro-Partitioning with up to 10 micro-partitions per processor, 254 per node; Multiple Shared Processor Pools; Virtual I/O Server; Shared Dedicated Capacity; PowerVM Lx86
PowerVM Enterprise Edition (optional)	PowerVM Standard Edition plus Live Partition Mobility
Operating systems	AIX V5.3 with the 5300-08 Technology Level SUSE Linux Enterprise Server (SLES) 10 SP2 for POWER or later; Red Hat Enterprise Linux 4.6 for POWER (RHEL4.6) or later; RHEL5.2 or later
Frame features	
Power requirements	200v to 240v; 380v to 415v; or 480v AC (3-phase at 50/60 Hz)
Cooling requirements	Chilled inlet water supply/return required for all systems.
System frame dimensions	79.5"H x 29.5"W x 60.0"D (201.3 cm x 75 cm x 152.4 cm); weight: 3,650 lb (1,656 kg)[2]
Warranty (limited)	9 hours per day, Monday through Friday (excluding holidays), next-business-day for one year at no additional cost; on-site (varies by country). Warranty service upgrades and maintenance are available.

For more information

To learn more about IBM Power 575 supercomputing nodes, contact your IBM marketing representative or IBM Business Partner, or visit the following Web sites:

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

[1] Based on IBM calculations of GFLOPS for a fully populated single Power 575 frame with 4.7 GHz POWER6 processors compared to a fully populated single System p5™ 575 frame with 1.9 GHz POWER5+ processors.

[2] Populated with 14 Power 575 nodes. Weight will vary when disks, adapters and other peripherals are installed.

[3] Based on IBM calculations of GFLOPS for 32 POWER6 processor cores running at 4.7 GHz (single node) compared to 16 POWER5+ processor cores running at 1.9 GHz (single node).

[4] Based on IBM calculations of energy use for Power 575 nodes versus a comparable air cooled system providing equivalent system performance.

IBM Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States
April 2008

POD03010-USEN-01