



IBM System p5 575
8-core 2.2 GHz, 16-core 1.9 GHz
Ultra-dense, Modular Cluster Nodes for
High Performance Computing

February 14, 2006

Harry M. Mathis
Frederick Bothwell
Jacob Thomas

IBM Systems and Technology Group

Introduction

IBM System p5™ 575 now applies IBM POWER5+™ chip technology in a cluster building block approach to help meet the requirements of high-performance technical and commercial computing. With the performance advantages provided by the powerful POWER5+ processors and their associated system architecture, a fast system bus, high memory bandwidth and robust input/output (I/O) subsystems, the 8-core and 16-core POWER5+ p5-575 systems provide the greatest sustainable memory bandwidth per fully populated rack of any IBM UNIX® environment clustered server. The p5-575 supports the AIX 5L™ V5.2, AIX 5L V5.3, SUSE Linux® Enterprise Server 9 (SLES 9) for POWER™, and Red Hat Enterprise Linux AS 3 (RHEL AS 3) for POWER or above operating systems (OSs). Its performance and enhanced clustering alternatives combined make the p5-575 a versatile solution for many demanding high performance computing (HPC) client requirements.



Figure 1: p5-575 Case Opened to Show CEC and Fans

The p5-575 is available both in 2.2 GHz 8-core or 1.9 GHz 16-core symmetric multiprocessing (SMP) server cluster nodes. Multiple p5-575 servers can be configured in distributed clusters. Up to 12 p5-575 modules can be contained in a single 24-inch IBM system frame, providing leading-edge performance in a very dense package. Building on over a decade of IBM RS/6000® SP™ and IBM @server® pSeries® 655 supercomputing experience, the POWER5+ p5-575, the latest development in the POWER5 p5-575 line, is ideal for computationally-intensive and data-intensive workloads in science, engineering, business intelligence and data warehousing. This white paper describes the newly available 8-core and 16-core POWER5+ p5-575 servers, including potential performance enhancing characteristics, extreme scalability and configuration versatility.

The p5-575, shown in Figure 1, reflects innovative, simplified modular packaging which is designed for enhanced reliability and extraordinary versatility. The p5-575 has eight memory DIMM slots

connected to the integrated memory controllers of each of the eight POWER5+ chips. The total of 64 DIMM slots supports an aggregate DDR2 memory capacity of up to 256GB per node, with a theoretical peak memory transfer speed of up to 25.58 GB/sec per POWER5+ chip (i.e., 25.58 GB/sec per core in the 8-core configuration or 12.7 GB/sec per core in the 16-core configuration). The DIMMs are located in close proximity to the associated POWER5+ chip, effectively minimizing signal propagation delays. In addition, the p5-575 implementation allows processors to communicate over a high-speed SMP fabric, as opposed to the I/O-based switch fabric required to interconnect multiple 2-core or 4-core systems. The larger SMP configuration of the p5-575 helps reduce the overall cost of computing by minimizing network infrastructure cost, base I/O and storage hardware cost and, in some cases, memory cost. In addition, fewer servers require less management and maintenance.

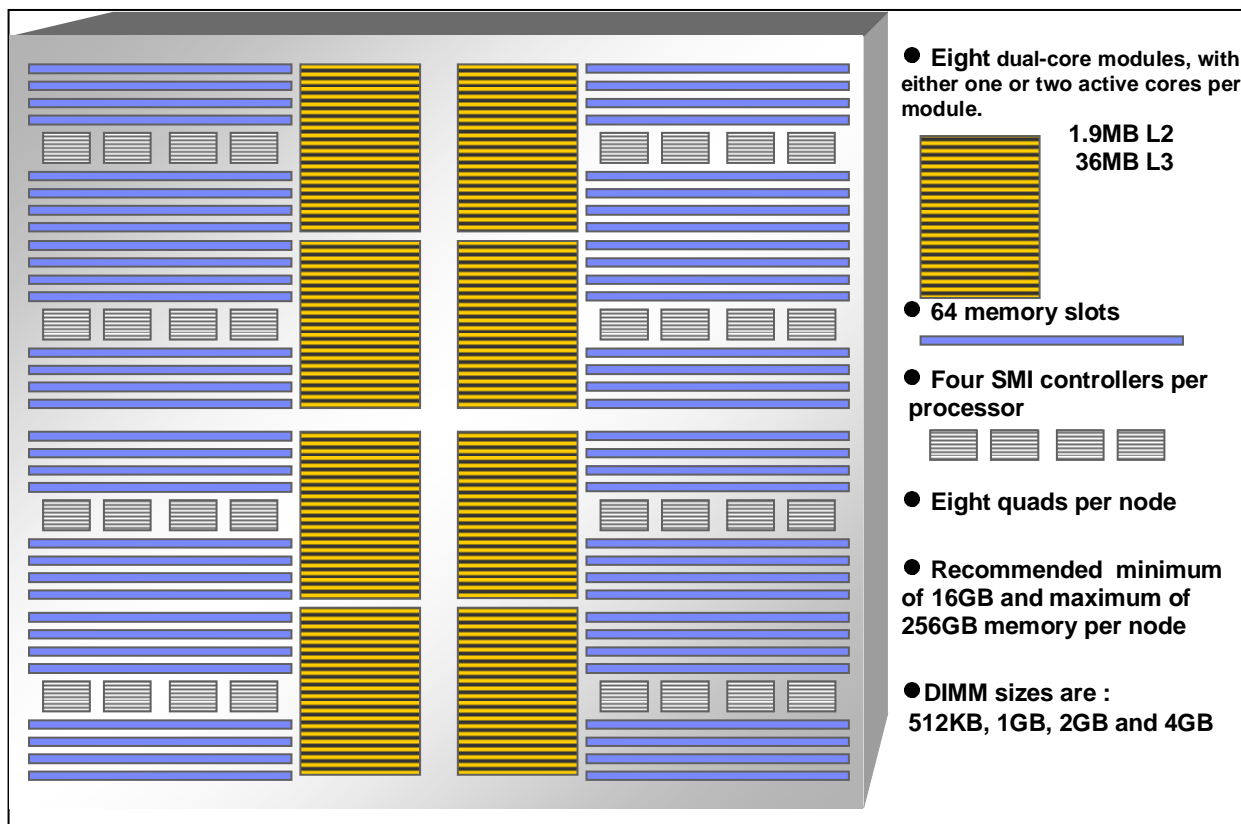


Figure 2: p5-575 CEC and Memory Subsystem

The p5-575 power distribution and conversion system – adopted from the IBM @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 designed 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 required to personalize for power utility voltage or frequency. Support for 200v to 240v; 380v to 415v; and 480v, 3 phase power inputs allow US clients to enjoy reduced facility equipment cost and help improve energy efficiency. The ability for the p5-575 to tolerate client utility power disturbances is exceptional in comparison to most other computing equipment, and optional internal battery back-up helps the p5-575 system ride through a momentary utility power interruption without the need for large and expensive UPS systems.

Finally, the unique p5-575 node implementation has only four field-replaceable components; the I/O subsystem, the DC power converter/lid, the processor and memory planar, and the cooling system. Each of these components are custom-designed to satisfy the demanding requirements of very high-performance, very high density computing.

The p5-575 as a Cluster Server

The p5-575 delivers exciting power, density and flexibility in a 2U by 24-inch wide by 48-inch deep package. The heart of the rack-mounted p5-575 cluster node is its processor/memory module, now available with an array of eight POWER5+ chips on a planar. In the 8-core configuration, each of the POWER5+ processor/memory modules contains a dual-core processor chip with only one core active and operating at 2.2 GHz, eight memory DIMM slots, and a private, high-performance, custom 36MB Level 3 (L3) cache chip. In the 16-core configuration, each of the POWER5+ processor/memory modules contains a dual-core processor chip with both cores active and operating at 1.9 GHz, eight memory DIMM slots and a private, high-performance, custom 36MB Level 3 (L3) cache chip. Each dual-core processor chip for the 8-core server contains one active processor core, a dedicated 1.9MB Level 2 (L2) cache, memory controller and L3 cache directory. Each processor chip for the 16-core server contains two active processor cores, a shared 1.9MB Level 2 cache, memory controller and L3 cache directory. Although all processor chip on-board-facilities operate at the processor chip frequency, 2.2 GHz for the 8-core system or 1.9 GHz for the 16-core system, the memory controller is designed to operate asynchronously, with the memory interface running at 1066 MHz to connect with the 533 MHz DDR2 memory.

In both the 8-core and 16-core configurations, the p5-575 server supports up to 256 gigabytes (GB) of DDR2 memory, over 300 megabytes (MB) of L2 and L3 cache, and provides a sustained memory bandwidth of 105.5 GB/sec. Also in both implementations a set of fabric buses operating at ½ the processor core frequency connect the eight POWER5+ chips to provide cache coherence, shared memory, and I/O functions.

The p5-575 memory subsystem consists of Level 1 (L1), L2, L3 caches and main memory. The L1 instruction cache and data cache are included within each processor core. The private 36MB L3 cache is located out of the path to main memory and operates at one-half of the chip frequency. Each processor is able to read from the L2 and L3 caches of the other chips, but can only store into its own L2 and L3 caches. Each p5-575 node contains a total of 15.2MB of L2 cache and 288MB of L3 cache.

The p5-575 provides 64 slots for DDR2 memory DIMMs. Each processor chip has point-to-point access to eight memory DIMMs divided into two sets of four DIMMs or quads. Memory for the p5-575 is available in 512MB, 1GB, 2GB and 4GB DIMMs and can be installed in the configurations listed in Table 1 and Table 2. Not shown in the tables below, in keeping with industry conventions, the p5-575 is orderable with a minimum memory of 2GB--one quad of 512 MB DIMMS--although IBM suggests at least 2GB per processor for most systems.

Case 1: Maximum Bandwidth									
		Fully Populated / Homogeneous							
Total Memory	GB Memory	Quads							
		1	2	3	4	5	6	7	8
(512MB DIMMs)	16	2	2	2	2	2	2	2	2
32	16	2	2	2	2	2	2	2	2
(1GB DIMMs)	32	4	4	4	4	4	4	4	4
64	32	4	4	4	4	4	4	4	4
(2GB DIMMs)	64	8	8	8	8	8	8	8	8
128	64	8	8	8	8	8	8	8	8
(4GB DIMMs)	128	16	16	16	16	16	16	16	16
256	128	16	16	16	16	16	16	16	16

Table 1: Memory Configurations for Maximum Bandwidth

Case 2: Minimum Memory Configurations									
		Half Populated							
Total Memory	GB Memory	Quads							
		1	2	3	4	5	6	7	8
(512MB DIMMs)	16	2	2	2	2	2	2	2	2
16	0	0	0	0	0	0	0	0	0
(1GB DIMMs)	32	4	4	4	4	4	4	4	4
32	0	0	0	0	0	0	0	0	0
(2GB DIMMs)	64	8	8	8	8	8	8	8	8
64	0	0	0	0	0	0	0	0	0
(4GB DIMMs)	128	16	16	16	16	16	16	16	16
128	0	0	0	0	0	0	0	0	0

Table 2: Minimum Memory Configurations

Although the p5-575 can be configured with 32 DIMMs, its best performance is obtained when all 64 slots are filled with DIMMs of the same size, as shown in Table 1. The minimum memory configurations, where each processor chip has half of its DIMM slots filled, are shown in Table 2.

Peak bi-directional bandwidth between each POWER5+ chip and its L3 cache is 30.4 GB/sec, or 243.2 GB/sec for the p5-575. Peak bi-directional bandwidth between each chip's memory controller and main memory is 12.4 GB/sec, or 99.7 GB/sec at the node level.

The p5-575 cluster node has two hot-swappable SCSI disk drives located at the rear of the enclosure. Both drives must be identical and are available in 36.4GB, 73.4GB, 146.8GB and 300GB sizes. The 36.4GB disk drive is available as a 15K rpm version; the 73.4GB and 146.8 GB disk drives are available as either 10K or 15K rpm versions; the 300GB disk drive is available as a 10K rpm version. Two dual 10/100/1000Mbps Ethernet ports and two ports connecting to the Bulk Power Hub Ethernet

switches elsewhere in the frame are located at the rear of the module. As an option, four hot-pluggable, blind-swap 64-bit, 133 MHz long PCI-X slots are available, along with the remote-I/O (RIO-2) port, and are located at the rear of the p5-575 enclosure, as shown in Figure 3.

Rear view of p5-575 node With PCI-X and RIO-2




Figure 3: Rear View of p5-575 Showing I/O Facilities

The RIO-2 hub port on the p5-575 node has the PCI-X option to support one I/O drawer. Two cluster nodes can share one 24-inch I/O drawer, which provides 20 hot-plug/blind-swap PCI-X slots in two 10-slot I/O domains, and contains up to 16 hot-swappable 36.4GB, 73.4GB or 146.8GB 15K rpm disk drives on four fully independent Ultra3 SCSI buses – for a maximum of 2.34 terabytes (TB) of disk storage per I/O drawer.

Major productivity enhancements are provided through the POWER Hypervisor in conjunction with the supported operating systems. The user can establish dynamic logical partitions (LPAR) running AIX 5L V5.2, AIX 5L V5.3, SLES 9 or RHEL AS 4 operating systems. Dynamic LPAR enables system administrators to reallocate system resources without rebooting the system or the partition.

If AIX 5L V5.3, SLES or RHEL AS are selected for a partition, the user can take advantage of the benefits of hardware simultaneous multithreading¹, which may provide an increase of up to 30% (based on rPerf projections²) processor throughput over single-threaded operation, depending on the nature of the applications being run in the partition.

Furthermore, with the AIX 5L V5.3 and Linux operating systems, the user can obtain even more flexibility with the optionally available Advanced POWER Virtualization feature, which provides Partition Load Manager (AIX 5L only), Micro-Partitioning™, shared processor pool and Virtual I/O Server capabilities. Partition Load Manager provides policy-based, automatic partition resource tuning that can adjust CPU and memory allocations in response to load demands. Micro-Partitioning provides the capability to establish up to 160 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 assigned to the pool to improve utilization and throughput, and which can be changed dynamically to meet changing environments. Virtual I/O Server 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.

The p5-575 Clustered System

IBM @server Cluster 1600 is a highly scalable cluster solution for AIX 5L or Linux operating systems, which include the AIX 5L V5.2 and above, SUSE Linux Enterprise Server 9 and Red Hat Enterprise Linux AS 3 and above 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) – a high function development and execution environment for parallel message-passing applications. PE is supported on AIX 5L.
- LoadLeveler® (LL) – dynamic job scheduling and workload balancing software supporting thousands of jobs within the cluster. LL is supported on AIX 5L and SLES.
- General Parallel File System (GPFS) – a high-performance, shared disk file system providing fast data access to all nodes in a cluster. GPFS is supported on AIX 5L and SLES.
- Engineering Scientific Subroutine Library (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 and SLES, while ESSL is supported on AIX 5L, SLES and RHEL AS.
- 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.

Since the p5-575 is a 2U high drawer-module system, 12 p5-575 servers along with two optional I/O drawers can be installed in a system frame (FC 5793; see Figure 4), which is 32-inches wide by 56-inches deep and has a footprint of 12.4 square feet. A Cluster 1600 system consisting of one system frame fully populated with 8-core p5-575 servers has 96 processors, and populated with 16-core p5-575 servers has 192 processors. Up to 128 p5-575 servers and 256 logical partitions (LPARs) can be interconnected in a Cluster 1600 system. Up to 128 individual p5-575 servers each with one or more LPARs until the maximum of 256 LPARs is reached can be clustered together to support large-scale computational modeling and multi-terabyte parallel databases.

The Hardware Management Console (HMC), required for LPARs and CSM-managed clusters, provides redundant network attachment to the p5-575 system frame to manage system control software, manage LPARs and control clusters, as well as to serve as the service focal point.

The p5-575 supports the Linux OS in either a SMP configuration (single partition mode) or in all LPARs provided a System p5, @server p5 or pSeries system running AIX 5L is network attached. The network-attached AIX 5L OS server is required for diagnostics and service support.

CSM is a cluster software tool designed to provide a foundation on which to scale-out hundreds of Cluster 1600 nodes and servers including the p5-575. Designed to deliver high performance and extreme horizontal and vertical scalability, CSM offers a highly effective clustered systems management; easy, continuous upgrades for growing workload requirements; and high levels of

system, data and application availability. The tools assist in hardware and software configuration and installation, device management, security administration, error logging, problem management, system recovery and resource accounting – all from a single point-of-control.

p5-575 24-inch Frame Node Layout




Figure 4: Clustered System With 12 p5-575 Servers and Two I/O Drawers

How p5-575 nodes are interconnected in a cluster is dependent on the cluster management software being used. The p5-575 is designed to support an industry standard gigabit Ethernet interconnection, the pSeries High Performance Switch (HPS) or 4x InfiniBand PCI adapters for Linux HPC workloads running TOPSPIN MPI (MVAPICH). Cluster Systems Management (CSM) supports Ethernet (10/100/1000 Mbps) and IBM HPS support for AIX 5L clusters and both 4x InfiniBand and Ethernet interconnection for Linux clusters.

The HPS is based on the proven technology and architecture of the SP Switch2 and, of the two available connectivity approaches, provides significantly greater communication bandwidth and lower latency for p5-575 nodes or their LPARs in Cluster 1600 environments. The HPS, a 4U rack drawer, fits in the 1U or 5U positions of the 24-inch frame. It provides a unified switch network with parallel, interconnected communications channels and supports copper interfaces for switch-to-switch connections. Redundant power converters and power cabling are designed to provide improved reliability, availability and serviceability (RAS). Up to 128 p5-575 cluster nodes and 256 switch-to-switch links are currently supported. Even higher scalability limits are available via special order.

The 42U system frame (FC 5793) supports the optional I/O drawer, as well as an optional battery backup subsystem. This I/O drawer adds 16 additional hot-swappable disk bays capable of housing 36.4GB, 73.4GB or 146.8GB 15K rpm Ultra3 SCSI disk drives. Thus a single p5-575 with I/O drawer can have 2.94 terabytes of disk storage. In addition, the I/O drawer contains 20 hot-

plug/blind-swap PCI-X bus slots supporting 64-bit adapters and offering backward compatibility for 32-bit cards. Two p5-575 modules are able to share one I/O drawer. I/O drawers must be placed in the same frames as the p5-575 servers to which they are attached. Thus a single system frame with five p5-575 servers and five I/O drawers installed may contain 14.74 terabytes of disk storage.

High Performance, High Density

The p5-575 server has achieved top ranking in a number of industry standard and application benchmarks. The p5-575 has achieved the highest SPECfp_rate2000 measurement and the highest SPECCompM2001 measurement for any 8-core server. Furthermore, both the 8-core and 16-core p5-575 systems have higher memory bandwidths than any other 8-core and 16-core high-density RISC-based systems respectively.

Benchmark	8-core 2.2 GHz p5-575	16-core 1.9 GHz p5-575
SPECint_rate2000	200	314*
SPECfp_rate2000	382	571*
SPECCompM2001 (chips/cores/threads)	40,560 (8/8/16)	56,211 (16/16/32)
LINPACK HPC	66,400 MFLOP/sec	111,400 MFLOP/sec
Stream Triad Standard	96,327 MB/sec	86,062 MB/sec
Stream Triad Tuned	100,523 MB/sec	86,379 MB/sec
MM5 V3.6 T3A Standard ³ (cores/threads)	14.89 (8/16)	N/A
MM5 V3.6 T3A Tuned ³ (cores/threads)	19.83 (8/16)	N/A

Table 3: Current as of February 7, 2006. Sources: <http://www.spec.org>,
<http://www.cs.virginia.edu/stream/>,
<http://www.netlib.org/benchmark/performance.pdf>,
<http://www.mmm.ucar.edu/mm5/mpp/helpdesk/20040304a.html>

* Submitted for publication February 14, 2006




Figure 5: 8-processor p5-575 Fluent Results⁴

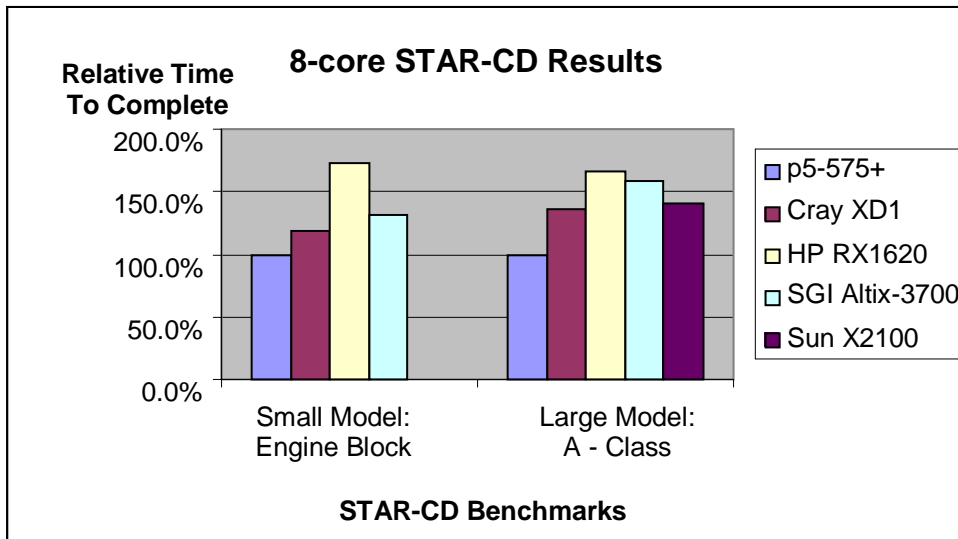


Figure 6: 8-processor STAR-CD Results⁵

The p5-575's outstanding performance results from several factors:

1. IBM's powerful POWER5+ processor
2. IBM's multi-chip packaging technology, which connects the POWER5+ processors with a wide, high-speed fabric supporting 204.6 GB/sec of peak memory to processor bandwidth
3. Integrated I/O slots which provide greater I/O bandwidth than blade server slots
4. Additional RIO-2 slots that provide enterprise size I/O capabilities and flexibility

Built-in Reliability Features

IBM autonomic computing enhancements have been architected into the p5-575. Self-protecting helps enable the p5-575 to determine the cause of an error as it happens and obviates lengthy service times attempting to recreate errors after the fact. Most errors are expected to be self-corrected or

functions varied off-line while the server remains available for use. IBM's exclusive First Failure Data Capture provides error information in real-time and makes it possible to determine the parts needed to fix the problem. The server has the capability to determine which part or component needs repair and phone IBM Service to identify parts needed for maintenance at a time acceptable to the client.

Self-healing capabilities help the p5-575 overcome error conditions and continue operating if a failure is detected. It 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 stores bits across four different memory words, enables recovery of single-bit errors and helps keep a p5-575 running when a failure is detected by Chipkill™ memory. Bit-steering dynamically routes a bit to a redundant memory chip in the event the memory failure rate for the bit exceeds a given threshold. If all bits are 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.

Exceptional system availability is facilitated through use of redundant power distribution to each node, hot-plug/blind-swap PCI and PCI-X slots, hot-swappable disks in the cluster node and optional I/O drawer, Dynamic Processor Deallocation in the event error thresholds are exceeded for a processor, and by the capability to vary individual processors of a p5-575 system off-line for maintenance.

Recommended Workloads

Recommended workloads for the p5-575 include high performance computing applications that are floating-point intensive and/or require high memory bandwidth. Typical HPC applications include: computational fluid dynamics, petroleum reservoir modeling, automotive noise/vibration/harshness analysis, atmospheric and ocean modeling, quantum chemistry, genomic sequencing and molecular modeling. The p5-575 is ideal for general technical workload consolidation, with the 8-core systems providing the maximum single-thread performance, cache size per core, memory bandwidth and capacity per core, and I/O bandwidth and capacity per core, while the 16-core systems provide more than 55% greater peak aggregate compute capacity integrated into the same cache, memory, and I/O infrastructure.

POWER5+ technology's support for simultaneous multithreading provides an additional level of flexibility, with each core capable of simultaneously executing two completely independent processes or threads. Simultaneous multithreading has been shown to provide significant increases in system throughput for a wide variety of scientific/technical/engineering (as well as commercial) workloads. Because simultaneous multithreading is not beneficial for all workloads, the p5-575 systems can be configured to run in either single-threaded or multi-threaded mode⁶.

In addition, large or small clusters of p5-575 systems are well suited for business intelligence applications such as data mining – using large, parallel databases that require a high degree of system scalability and dense packaging. Since the Cluster 1600 software enables up to 128 p5-575 cluster nodes to be integrated with other System p5, @server p5 and pSeries systems, it is a good way for clients to upgrade an existing pSeries 655 system or to integrate more computing power.

The p5-575 is also an excellent choice where floor space is at a premium and density is a significant advantage.

Optional Rear Door Heat eXchanger™ (RDHx)

A new technology component that uses the existing chilled water supply from installation air conditioning systems is now available to reduce System p5 server heat emissions by up to 55 percent. IBM was the first systems vendor to develop such a technology, introduced last year for deployment with IBM @server xSeries™ system. Now available for System p5 19-inch racks and 24-inch frames, RDHx eases the burden on existing air conditioning units and potentially lowers energy costs by up to 15 percent.

Designed to replace the rear door of the p5-575 24-inch frame, the Heat eXchanger employs sealed tubes filled with circulating chilled water to remove up to 55 percent of the heat generated in a fully populated rack and dissipate it so it is not released into the datacenter. The Heat eXchanger can remove up to 50,000 BTU of heat generated by a full server rack, based on total rack output.

Its unique design uses standard fittings and couplings with no moving or electrical parts. It can be opened like any rear cover, so serviceability of racks or frames fitted with a Heat eXchanger is as simple as a standard air cooled rack or frame.

Summary

IBM has been a leader in designing, developing and implementing clustered systems for over 10 years. The p5-575 is the latest addition to the Cluster 1600 family and the logical successor to SP clusters where the goal is high performance and extreme scalability for computational and data-intensive workloads in technical computing or business intelligence. It offers leading-edge POWER5+ technology and along with the p5-595, is a cluster building-block for high-performance systems. Cluster 1600 software enables clients to set up clustered p5-575 systems or seamlessly add the p5-575 server to an existing System p5, @server p5 and pSeries cluster, thereby enhancing and helping to protect their investment.

The p5-575 provides leading-edge performance, proven technology, excellent density and application and tool availability, along with the high-quality support and reliability features that IBM clients are accustomed to expect. This cluster-optimized, high-density server is the ideal choice for scientific and engineering research centers and data centers supporting high-priority technical and business intelligence workloads.

Footnotes

¹ Not supported on AIX 5L V5.2

² 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. 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 system 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.

³ The PSU/NCAR mesoscale model (known as MM5) is a limited-area, non-hydrostatic, terrain-following sigma-coordinate model designed to simulate or predict mesoscale atmospheric circulation. The model is supported by several pre- and post-processing programs, which are referred to collectively as the MM5 modeling system. The MM5 modeling system software is mostly written in Fortran, and has been developed at Penn State and NCAR as a community mesoscale model with contributions from users worldwide. See <http://www.mmm.ucar.edu/mm5/> for further details.

⁴ FLUENT is a leading Computational Fluid Dynamics (CFD) Software package. FLUENT's software is used for simulation, visualization, and prediction of fluid flow, heat and mass transfer, and chemical reactions. FLUENT along with its suite of pre- and post-processor software solutions covers the entire CFD modeling process – concept, design, analysis and simulation. See <http://www.fluent.com/software/fluent/index.html> for further details.

- FL5S1: 32,000 hexahedral cells: Turbulent flow in a bend, segregated implicit solver
- FL5S2: 32,000 hexahedral cells: Turbulent flow in a bend, coupled implicit solver
- FL5S3: 89,856 hexahedral cells: Flow in a compressor, Rotor 37, coupled implicit solver
- FL5M1: 155,188 tetrahedral cells: Coal combustion in a boiler, particle tracking, segregated implicit solver
- FL5M2: 242,782 hybrid cell: Turbulent flow in an engine valve port
- FL5M3: 352,800 hexahedral cells: Combustion in a high velocity burner
- FL5L1: 847,746 hexahedral cells: Transonic flow around a fighter, coupled implicit solver
- FL5L2: 3.6 Million hybrid cells: External Aerodynamics Around a Car Body, segregated implicit solver
- FL5L3: 9.8 Million hexahedral cells: Turbulent flow in a transition duct, segregated implicit solver

⁵ STAR-CD is a general purpose, unstructured grid, finite volume CFD code. It is efficiently programmed and highly parallel using the MPI communication protocol. STAR-CD is able to solve steady state or time dependent problems with the optional solution of moving grids, thermal transport, chemical reactions, multiphase fluids, and species concentrations. STAR-CD is used in many fields including the automotive industry, aerospace, turbomachinery, biosciences, building ventilation and safety, and combustion processes. See <http://www.cd-adapco.com/products/STAR-CD/performance/320/index.html> for further details.

⁶ Details of the control of single-thread and multi-thread mode depend on the specific levels of software installed.