IBM System Blue Gene Solution Enables Breakthrough Science for Healthcare and Life Sciences

Challenges in the Health Care and Life Sciences Demand Leadership Computing Power

In the Healthcare and Life Sciences (HCLS) industry, the focus is on improving efficiency, productivity and performance to reduce development timelines and get through clinical trials sooner. Key to success is the ability to apply computational power to increasingly complex and demanding research and business processes.

Workloads such as chemical and biological simulations, medical imaging and genome assembly can all benefit from the application of greater computational power.

Today, increasingly complex models and growing amounts of data are intensifying the need for increased memory and compute power, now more than ever. Data complexity is growing faster than it can be absorbed with traditional methods. And, it is becoming more common for important jobs to run for ever longer periods of time, impacting competitiveness.

Up to now, some organizations would be forced to accept compute limitations on their research and development due to financial constraints. Others would try to assign more processors to a job, hoping that additional servers would not produce too much heat, require too much floor space, and consume too much power. Clearly, a new approach is needed that can offer high performance and extreme scalability in an efficient, affordable package that provides a familiar environment to the user community.

The Blue Gene System is Designed to Meet these Challenges

The IBM® System Blue Gene® Solution is the result of an IBM supercomputing project which began in collaboration with Lawrence Livermore National Laboratory over five years ago. It was dedicated to building a new and innovative family of supercomputers optimized for bandwidth, scalability and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today's high performance systems. The level of performance provided by the Blue Gene system can enable a tremendous increase in the scale of biological simulations beyond what is possible with other supercomputers. Hands-on experiences with a variety of applications have shown that the Blue Gene architecture is applicable to a number of parallel workloads found within HCLS and Life Sciences disciplines.

The Blue Gene system is not just a supercomputer that delivers ultrascalable performance. It is also extremely efficient. Because of unique design points that allow dense packaging of processors, memory and interconnect, the Blue Gene system offers leadership efficiency in floor space and power consumption.

The Blue Gene system enables breakthrough science, handling problems that cannot be solved on traditional platforms, offering clients new insights and providing them a tremendous competitive advantage.

Available in configurations ranging from one to 64 racks, The Blue Gene system is the innovative new solution from IBM that delivers an ultrascalable solution without sacrificing efficiency.

Innovated to Enable Breakthrough Science

The Blue Gene system is built out of a very large number of compute nodes, each of which has a relatively modest clock rate contributing to both low power consumption and low cost. The Blue Gene system utilizes IBM PowerPC® embedded processors, embedded DRAM and system-on-a-chip techniques that allow for integration of all system functions including compute processor, communications processor, three cache levels, and multiple high speed interconnection networks with sophisticated routing onto a single chip. Because of a relatively modest processor cycle time, the memory is close, in terms of cycles, to the processor. This is also advantageous for power consumption and enables construction of dense packages in which 1024 dual-processor compute nodes can be placed within a single rack.

The nodes are interconnected through five networks: a 3-dimensional torus network for point-to-point messaging between compute nodes, a global collective network for collective operations over the entire application, a global barrier and interrupt network, and two gigabit Ethernet networks for machine control, and for connection to other systems. The torus network is particularly effective for applications with locality of communication. And the global collective network is useful for speeding up MPI collective communications constructs.

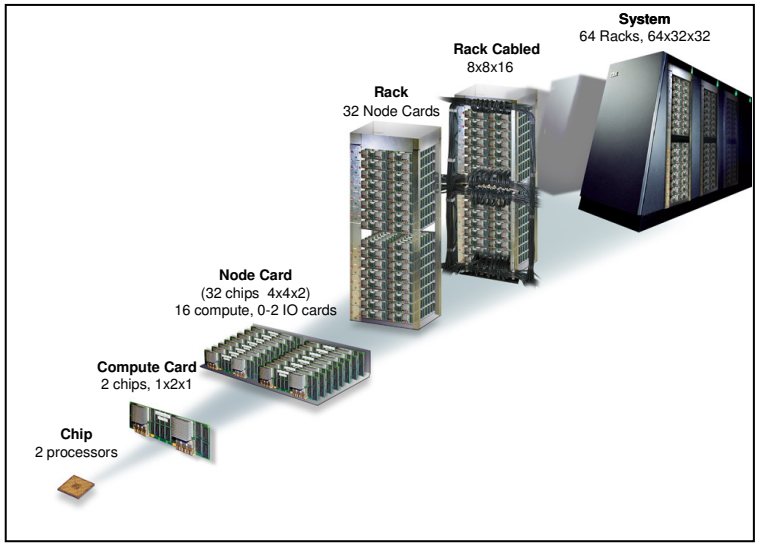

A Familiar, Tuned Software Environment

Three fundamental principles were followed when the system software was designed for the Blue Gene system: simplicity, performance and familiarity. Driving toward simplicity in the software design has allowed development of software that takes advantage of hardware features to deliver high performance without compromising stability and security. And by creating a programming and administration environment based on familiar programming languages, libraries, job management tools and parallel file systems, clients benefit from the innovative design elements of the Blue Gene system without facing a steep learning curve.

The Blue Gene System is Available On Demand

IBM provides access to the Blue Gene system through the Deep Computing Capacity on Demand (DCCoD) center. Clients with constrained budgets and limited need for access can request time on the system and just pay for the amount of capacity reserved subject to availability. In this way, clients can contract for variable capacity and services to help satisfy short term planned or unplanned peak workloads. Subject to availability and scheduling, accessing the Blue Gene system through the DCCoD center can help clients quickly tap supercomputing power while helping to reduce financial and technical risk.

Backed by IBM

The Blue Gene system is backed by a one-year maintenance package that covers all hardware and software components. Support is provided Monday through Friday 8 A.M. to 5 P.M. local time from the local IBM support team. An extended maintenance contract is available which provides parts replacement and software support for years following the first year of coverage.

Ultrascalability Benefits Genome Assembly

By reducing computation time from hours to minutes to even seconds, the Blue Gene system offers a new class of systems capability. A single rack has a peak performance rating of 5.7 trillion operations per second.

The Blue Gene system is an excellent platform for running extremely large parallel workloads by virtue of its ultrascalability. Genome assembly is one example.

Genome assembly, which is used for detection of diseases, development of therapies, and improving crop yield involves analysis of tens of millions of short genomic fragments. Genome assembly is an ideal application for parallel processing, but requires huge computational power, large amounts of memory and efficient internode communication. The effort becomes even more demanding when a large portion of the genomic fragments are repeats of other fragments.

The Blue Gene system's massively parallel capability and special-purpose interconnects allow it to perform well on genome assembly problems. Efficiency is maintained when increasing the number of processors while increasing the size of the problem. This allows users to handle large, complex, gene-enriched models beyond what has been possible with other systems.

Users also have flexibility in how the Blue Gene system's dual-processor nodes are deployed. Both processors may be used for computation, or one may be used for computation while the other can be dedicated to communication. This latter approach, known as "co-processor mode" is valuable for workloads like genome assembly where internode communication traffic is high.

Extending the Length of Molecular Dynamics Simulations

Classical molecular dynamics involves the modeling of interactions between particles. In the case of biomolecular simulation, special care must be taken with regard to long-range electrostatic interactions to ensure behavior is properly modeled.

To gain insight, researchers continually need to drive up the simulation time. Running molecular dynamics on large-scale parallel machines is the most promising way to extend the length of simulations to the microsecond scale and beyond.

or

There are two main challenges to running such workloads in parallel: achieving a balanced distribution of work across large numbers of compute nodes, and ensuring internode communication is efficient. The Blue Gene system is supported by sophisticated software that enables very fine load balancing so that thousands of nodes work in harmony. Also, the Blue Gene system's special-purpose networks are designed for highly efficient communication between and among nodes. As such communication is often the limiting factor in runtimes of molecular dynamics simulations on parallel systems, superior networking capability is critical.

Power and Floor Space Efficiency

Today, many customers are reaching the limits of the power and floor space available in their computer complexes, and they are finding it difficult or impossible to deal with the heat dissipated by densely packed servers based on commodity chip technology.

The Blue Gene system was specifically designed to deliver the most performance per kilowatt of power consumed, and per square foot of floor space occupied. Blue Gene is a leader among supercomputers is more power and space efficiency.

The key to maximizing performance per rack is to maximize processor performance per watt. The low-frequency, low-power embedded processors in the Blue Gene system are superior to today's high-frequency, high-power microprocessors by a factor of 2 or more. This design principle – using large numbers of low-power system-on-a-chip technology – allows the Blue Gene system to outperform commodity clusters while saving on power and floor space.

A Commitment to Deep Computing

IBM Blue Gene system was developed to enable breakthrough science with innovations that result in leadership performance, ultrascalability, and environmental efficiencies. The Blue Gene system is accompanied by a product roadmap that stretches to petaflop performance.

The Blue Gene system can bring significant competitive advantage to Healthcare and Life Sciences organizations that require solutions to their most demanding, computationally intensive workloads.

For more information

To learn more about the IBM System Blue Gene Solution, please contact your IBM marketing representative or visit the following Web site: **ibm.com**/servers/deepcomputing/bluegene.html

Details At a Glance

Attribute	Details	Benefits
Processor	PowerPC 440 700MHz; two per node	Low power allows dense packaging; better processor-memory balance
Memory per node	512 MB SDRAM-DDR (Model 0203-700) 1 GB SDRAM-DDR (Model 0203-900)	
Networks	1) 3D Torus - 175MB/sec in each direction 2) Collective Network – 350MB/sec; 1.5 usec latency 3) Global Barrier/Interrupt 4) Gigabit Ethernet (I/O & connectivity) 5) Control (system boot, debug, monitoring)	Special networks speed up internode communications; designed for MPI programming constructs; improve systems management
Compute Nodes	Dual processor; 1024 per rack	Double FPU improves performance
I/O Nodes	Dual processor; 16-128 per rack	Facilitates job launch and I/O, raising efficiency of compute nodes
Operating Systems	Compute Node – Lightweight proprietary kernel I/O Node – Mini Control Program Front End and Service Nodes – SuSE SLES 9 Linux	Kernel tailored to processor design; industry-standard distribution on front-end and service nodes preserves familiarity to end users and administrators
Performance	Peak performance per rack – 5.73 TFlops Linpack performance per rack – 4.71 TFlops	Highest available performance benefits capability customers
Power	27.6 kW power consumption per rack (maximum) 7 kW power consumption per rack (idle) 208 VAC 3-phase; 100 amp service per rack	Low power draw enables dense packaging
Cooling	Air conditioning 8 tons/rack (minimum) 2800 CFM (compute rack); 350 CFM (power supplies)	Low cooling requirements enable extreme scale-up
Acoustics	9.0 LwAD and 8.7 LwAm	
Dimensions (includes air duct)	Height – 77" Width – 36" Depth – 36" Weight – 1810 lbs. Service clearances – 30" front and back Raised floor height – 16" minimum	Design allows dense floor plan layout for better floor space utilization

IBM Corporation
Integrated Marketing Communications
Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States
May 2006