IBM System Blue Gene Solution Provides Competitive Advantage for Upstream Petroleum Companies

Challenges in the Petroleum Industry Demand Leadership Computing Power

In the Petroleum Industry, the drive to increase reserves drives everything. Key to success is the ability to formulate, run and analyze increasingly complex computationally intensive simulations. Companies need to locate and optimize hydrocarbon reserves faster by making better decisions with lower risks. And that means they require solutions to help improve their competitive position in exploration and reservoir management.

In the area of seismic data processing, systems that offer high performance and extensive scalability are needed to improve oil and gas exploration and hydrocarbon recovery. Companies must have the capability to perform larger, more complex, time-sensitive analyses. They must be able to speed the process of first-time-to-oil. And companies continually strive to identify promising acquisition blocks and improve drilling success rates.

As for production, it is absolutely critical that companies manage reservoirs so that maximum product is derived in the most cost-effective way. Again they need tools to help build better and more accurate reservoir models. That means companies must increase efficiency with virtual real-time data capture and analysis. And they must employ parametric studies and compute-intensive algorithms that yield higher fidelity, lower-risk results.

Such computationally intensive simulation and analysis demands leadership computing power. Companies' competitiveness is directly related to the strength of their computing capabilities. IBM® System Blue Gene®, the world's fastest ultrascale supercomputer when configured with 64 racks enables breakthrough science that can deliver the competitive advantage required in the Upstream Petroleum field.

Designed to Meet the Challenges

The IBM Blue Gene Solution is the result of an IBM super-computing 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 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 the Petroleum industry.

The Blue Gene system is not just a supercomputer that delivers ultrascalable performance. 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, providing clients new insights and competitive advantage.

Available in configurations 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 with 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.

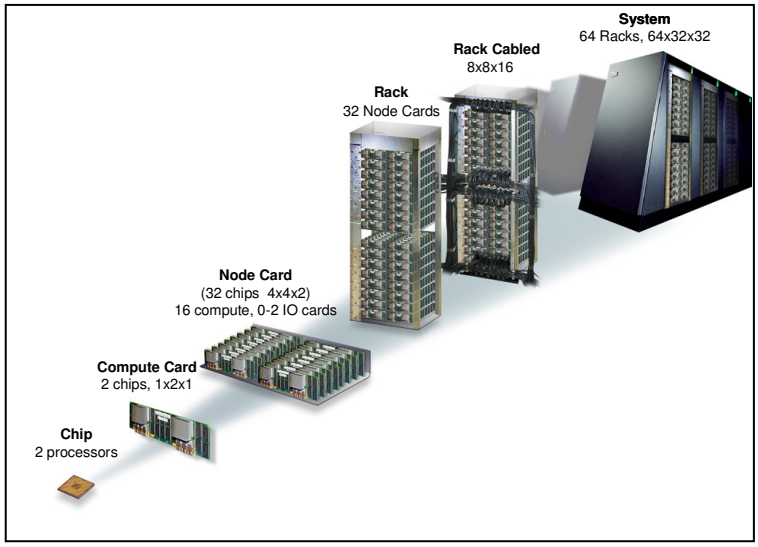

A Familiar, Tuned Software Environment

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 also 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 Seismic Data Processing

By reducing computation time from days to hours to even minutes, 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 seismic workloads by virtue of its ultrascalability.

New research is being conducted by universities and within the petroleum industry to exploit parallel computation to a greater degree on reflection seismic data processing. It is through reflection seismology that engineers can get the most detailed view of underground reservoirs for petroleum exploration and production (E&P).

One focus area today involves estimating seismic wave velocity from waveform data. To accomplish such computationally intensive techniques requires large-scale parallel capability. Experiments conducted jointly with major oil companies on their proprietary applications indicate that Blue Gene is very applicable to seismic data processing.

The Blue Gene system's massively parallel capability and special-purpose interconnects allow it to perform well on finite element methods, differential semblance optimization, and other simulation techniques. Efficiency is maintained when increasing the number of processors while increasing the size of the problem. This allows users to handle large, complex, models beyond what has been possible with other systems.

Petroleum Recovery Depends on Detailed Reservoir Modeling

At the heart of reservoir modeling is the characterization of the flow of substances within porous media. The solution involves computational techniques such as nonlinear partial differential equations, finite-difference and finite-element analysis, and domain decomposition methods.

The Blue Gene system excels when applied to these computational techniques because of its superior balance of floating point operations and internode communications. Users of the Blue Gene system will see more communications volume per floating point operation than on other systems, and this translates to faster solutions.

Extracting oil and gas in the most efficient, cost-productive way is essentially an optimization problem in which there is incomplete information about the characteristics of the reservoir. Traditionally, this problem has been a time-consuming part of the reservoir modeling process. Researchers are working on development of better parallel solvers, and the early results indicate that Blue Gene is a strong platform for optimizing oil and gas recovery.

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.

Blue Gene and Deep Computing Visualization: A Winning Combination

Finding the right drill location requires the skills of geologists, geophysicists, seismologists, and reservoir engineers—and each uses a different application. This produces large sets of complex, multi-dimensional, nonlinear data. And as primary research and data capture improves through the use of the Blue Gene system, the amount of data increases exponentially. IBM offers Deep Computing Visualization (DCV), a tool that helps engineers quickly visualize information to understand, analyze and mine the data to gain insights and make better decisions.

DCV enables the display of applications on large multi-projector display walls, caves or high resolution monitors with no sacrifice in quality or speed. The DCV solution includes commodity-based hardware and software tools for ease of implementation and upgrades, and open standards-based middleware to allow intensive graphics rendering in a centralized location and high performance transmission across geographies.

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 the Petroleum Industry that requires 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
June 2006