

IBM System Blue Gene Solution Delivers Competitive Advantage to Finance and Reinsurance Firms



Challenges in the Finance Industry Demand Leadership Computing Power

In the Finance Industry, the most competitive firms are those that can maximize returns while minimizing risk, all in the shortest time possible. Key to success is the ability to apply computational power to increasingly complex and demanding business processes.

Workloads such as risk management, portfolio analysis, derivatives pricing and actuarial simulations can all benefit from the application of greater computational power.

Today, swelling datasets are intensifying the need for increased memory and compute power, now more than ever before. Data complexity is

growing faster than it can be absorbed with traditional methods. And, it is becoming more common for overnight batch jobs to finish outside the allotted window of time, yielding less-valuable results.

Firms can no longer accept responding to market conditions a day later; financial engineers must be able to respond to real-world events as they occur.

Traditionally, more compute processors would simply be assigned to a job, but this is no longer a viable solution as additional servers produce too much heat, require too much floor space, and consume too much power. A new approach is needed that can offer high performance and extreme scalability in an efficient package that provides a familiar environment to the user community.

Blue Gene is Designed to Meet these Challenges

The IBM® System Blue Gene® Solution is the result of an IBM supercomputing project begun over five years ago dedicated to building a new 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 Blue Gene can enable a tremendous increase in the scale of financial 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 finance and securities disciplines.

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

Blue Gene, the world's fastest computer, is the kind of tool that enables breakthrough financial analysis. Blue Gene handles 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, Blue Gene is the innovative new solution from IBM that delivers an ultrascale solution without sacrificing efficiency.

Blue Gene: Innovated to Enable Breakthrough Financial Analysis


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.

Blue Gene 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.



Familiar Software Environment Tuned for Blue Gene

Three fundamental principles were followed when the system software was designed for Blue Gene: 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 Blue Gene without facing a steep learning curve.

Blue Gene is Available On Demand

IBM provides access to

Blue Gene 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 Blue Gene through the DCCoD center can help clients quickly tap Blue Gene supercomputing power while helping to reduce financial and technical risk.

Backed by IBM

Blue Gene 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 with a maximum two-hour response time objective 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.

Faster Simulations Mean More Timely Solutions

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

Blue Gene is an excellent platform for running extremely large simulations by virtue of its ultrascaleability. Monte Carlo techniques are one example.

Monte Carlo techniques can be applied to a variety of financial engineering problems such as determining stochastic volatility, and pricing of multidimensional options which can include index options and basket options.

Monte Carlo techniques involve the generation of random numbers that seed independent processes that simulate the distribution of underlying prices and rates of return. The mathematics is such that a 10x improvement in the accuracy of the results requires a 100-fold increase in the number of simulation runs. Clearly, supercomputers are an appropriate tool for these workloads. Because Monte Carlo simulations run well in parallel and scale nicely they are well-suited to Blue Gene.

A Paradigm Shift in Memory Management

Packing processors more densely than ever possible before, IBM's Blue Gene represents a paradigm shift in computing. Just imagine what

can be done with a terabyte of memory attached to 1,024 multiply-interconnected nodes. IBM's Blue Gene has altered the ratio of CPU to memory.

Distributed shared memory gives the flexibility and scalability needed for real-time services-oriented computing. Blue Gene maintains a constant ratio of dedicated I/O nodes to compute nodes, and allows near linear scaling on the I/O bandwidth to compute nodes. Business applications have been written for Blue Gene that work with a variety of data sources and formats including streaming data and database-resident data.

Blue Gene Provides Strategic Insight

Today, Blue Gene can analyze a five petabyte datastream in real-time to allow analysts to recognize new trends, identify affinities, extract hidden patterns and isolate market opportunities as they emerge.

Furthermore, financial institutions can no longer stay competitive by simply analyzing existing trends or "strong signals." Blue Gene makes it possible to identify "weak signals" – those which announce future changes in a given landscape. It enables businesses to identify ripples in an ocean of information that can be tracked and analyzed before developing into waves with market impact.

Blue Gene Operates Within Power and Floor Space Barriers

For many years, financial firms simply added servers and blades to their rack systems as demand increased for processing capability. This strategy is quickly coming to an end, as customers reach the limits of the power and floor space available in their computer complexes. Furthermore, customers are finding it difficult or impossible to deal with the heat dissipated by densely packed servers based on commodity chip technology.

Blue Gene was specifically designed to deliver the most performance per kilowatt of power consumed, and per square foot of floor space

occupied. No system is more power- and space-efficient than Blue Gene.

It has been established that the key to maximizing performance per rack is to maximize

processor performance per watt. The low-frequency, low-power embedded processors in Blue Gene 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 Blue Gene to outperform commodity clusters while saving on power and floor space.



or



A Commitment to Deep Computing

IBM Blue Gene was developed to enable breakthrough financial analysis with innovations that result in leadership performance, ultrascaleability, and environmental efficiencies. Blue Gene is accompanied by a product roadmap that stretches to petaflop performance before the end of the decade.

Blue Gene can bring significant competitive advantage to Financial firms 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

©Copyright IBM Corporation 2006

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

Produced in the United States
April 2006
All Rights Reserved

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features or services discussed in this publication in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

IBM, the IBM logo, Blue Gene, POWER and PowerPC are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

A full list of U.S. trademarks owned by IBM may be found at: ibm.com/legal/copytrade.shtml.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

Blue Gene 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 - Embedded Linux 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