Deep Computing Capacity on Demand: Financial Risk Management and Compliance in Capital Markets and Retail Banking

Highlights

- *Addresses peak workload demands for high performance computing*
- *Offers flexible pricing, allowing customers to pay only for the capacity reserved and helps avoid up-front capital investment*
- *Helps create a flexible, scalable and responsive IT infrastructure*

IBM Deep Computing

IBM Deep Computing can address the needs of customers who require powerful solutions to address their most complex and challenging business needs. These customers include scientists, engineers, knowledge workers, and other professionals whose core business or research relies on high performance computing (HPC) for deep levels of computation, data access, visualization, and communications. Deep Computing is more than traditional HPC hardware, software and services. It also embraces emerging strategic technologies, research initiatives, open standards, industry and customer partnerships, and industry expertise. IBM Deep Computing takes a comprehensive approach to helping our customers gain insight and transform their business and organizations. Industries and business segments include Life Sciences, Digital Media, Financial Services, Petroleum, Electronics, Automotive, Aerospace, Government and Higher Education.

Using HPC in the Financial Services industry

The risk management and compliance discipline of the financial services industry is well-suited to take advantage of an on demand computing model.

Financial service companies deploy high performance computing (HPC) infrastructures to support processes such as trading analytics, customer analytics and segmentation, fraud detection, anti-money laundering, transaction monitoring for insider trading, credit limit decisions, and compliance with statutes like the Sarbanes-Oxley Act, Basel II (see sidebar "The impact of Basel II on risk management") and the USA Patriot Act.

A financial institution's HPC infrastructure can provide a scalable, flexible, resilient and affordable IT environment that yields competitive advantage. Leveraging large-scale variable HPC capacity on demand when needed can help improve risk management and compliance application performance. More sophisticated and accurate analytics are possible with an increase in granularity and scale of modeling functions. The ability to run several risk scenarios in parallel to generate an optimal solution as quickly as possible, without sacrificing accuracy and performance, is an advantage.

Near real-time analysis is possible. Risk management and compliance application run times can be reduced from overnight to intra-day and from hours to minutes. And business decisions can be accelerated. Cost and time of reporting can be reduced and clients can comply with regulations more efficiently and cost effectively.

IBM's offerings for on demand delivery of high performance computing can help financial companies meet demands for peak capacity and fluctuating workloads, while providing a flexible approach for acquiring and deploying computing resources and managing capital and operating costs. Financial companies can achieve attractive price/performance for computational- and data-intensive

risk management and compliance applications. These companies may also reduce the administrative costs associated with managing distributed systems and data, allowing them to focus on their core business functions rather than on IT cost recovery or infrastructure complexity.

Meeting the challenges facing financial risk and compliance managers

Today's climate of increased economic and regulatory challenges is driving financial services companies to reevaluate their approaches to risk management and compliance. Looking beyond market risk, these firms are examining credit, liquidity, operational and other risks. For instance, the movement from market to credit risk greatly increases the complexity of computations because:

- *Longer time frames with more steps are involved; market risk typically has a 24 hour hold period, whereas credit risk can have as much as a 10 year hold period.*
- *Credit mitigants affect credit exposures, including netting and collateral agreements, and must be modeled as part of the exposure.*
- *Complex corporate relationships must be understood to examine the exposure to any given entity.*
- *Some decisions cannot be considered in isolation—a whole portfolio of decisions must be optimized.*
- *Decisions must not be automated at the expense of business agility—these decisions must be able to evolve and change over time.*

The complexity of products and portfolios requires more value-at-risk versus Monte Carlo calculations and scenario simulations. In addition, financial services risk managers need their analyses to move from overnight to intra-day and even to near real-time for incremental value-at-risk calculations, price equilibrium models, pre-deal credit limit checking and other functions. Companies have large volumes of historical data, but the data can be leveraged only if accurate models can be executed and deployed timely.

Regulations and recent examples of noncompliance are demanding that firms place significantly greater emphasis on compliance reporting for things like anti-money laundering and transaction monitoring for insider trading, fraud detection in retail banking

The impact of Basel II on risk management

Basel II will require firms to have a single risk architecture across the entire organization, thus presenting an integrated and auditable picture of a firm's risk. To successfully implement this architecture, financial institutions must incorporate financial data and the risk management process into the firm's complete business outlook. Current infrastructures and risk models must include credit, liquidity and operational risk. To obtain this integrated view of risk, firms must operationalize their approach for credit risk calculations in Basel II (for example, Standardized, Foundation internal ratings-based (IRB) or Advance IRB). In turn, by deploying integrated decision engines behind all applications they can better ensure consistent operational decision making.

This integration poses a significant challenge to completing comprehensive and accurate risk reporting across an organization. IBM's Deep Computing Capacity on Demand capability extends the financial institution's in-house HPC infrastructure. It provides a scalable, highly resilient, and highly secure on demand operating environment and allows firms to use their existing custom-built and best-of-breed commercial analytics applications and tools (including leading risk management and compliance software from Algorithmics, Fair Isaac, SunGuard, SAS, Moody's KMV, and many more). IBM's Deep Computing Capacity on Demand capability can help financial institutions efficiently integrate their risk management processes while improving analytical performance and regulatory compliance.

and risk compliance for statutes such as Sarbanes-Oxley and the Basel Capital Accord (Basel II). Effectively tracking and analyzing potential high-risk activities can be challenging, given the increasing complexity of financial products, vast amounts of real-time and historical data to mine, and larger trading and transaction volumes. To prevent abuses before they can occur, firms are looking for a more systematic and near real-time means for identifying patterns of activity.

The greater volume and complexity of data and desire to run multiple risk scenarios simultaneously has spurred an unrelenting demand for additional processing power and speed, stretching IT resources well beyond their current capabilities. Few firms can make the capital expenditures required to address these computing requirements and build sufficient HPC capacity to satisfy peak workloads—especially given the immense pressure on costs

these firms are facing today. Risk and compliance managers need real-time, accurate results without sacrificing cost, liquidity or competitiveness.

Examining HPC requirements for financial risk management and compliance

Providing HPC resources on demand requires an infrastructure of high-speed processors and storage that can provide measurable compute and data storage capacity in a flexible, scalable and highly secure manner. Users require seamless access to these resources, regardless of where the physical systems may reside. An HPC infrastructure for the financial services industry has the following objectives:

Scalability: Applications and databases often exceed the capacity of even the largest single system, so systems need to grow with demand.

Security: Risk and compliance managers need a security-rich processing and information environment.

Flexibility: While they focus on controlling costs, financial institutions need increased flexibility in a rapidly changing environment.

Significant processing capacity: As processing requirements continue to grow, more and more companies are finding that they can benefit from using clustered servers and Grid computing. Clusters link separate nodes to create a single, powerful, scalable system, while Grid computing increases overall compute resource sharing and utilization.

Flexible financial and delivery models: Financial institutions can benefit from balancing their fixed and variable IT infrastructure and operations costs and choosing between onsite owned/operated versus offsite hosted IT solutions.

Clustering and Grid computing, combined with a wide range of server and storage technologies, are designed to provide:

- *High availability with failover protection*
- *Scalability with minimal downtime*
- *Large-capacity data volumes*
- *Capability to handle peaks in workload, permitting flexible responses to changing business requirements*

Providing high performance computing on demand

Fundamentally, providing HPC capacity on demand is a simple concept. Instead of purchasing computers to meet peak demands, companies can procure sufficient hardware for average demand and then contract remote processing power to help meet peak loads, or offload all processing to a remote facility. In practice this process can be more

complex. However, IBM can help shield customers from much of this complexity—one of the primary advantages of a hosted on demand environment.

High performance computing can have many different requirements, and purchasing sufficient hardware can become expensive. Some problems are computationally parallel, and a large number of servers working in parallel can solve the problem quickly and efficiently. Other problems require ultrafast interprocess communication or a single large memory model. A hosted environment can supply access to these resources, and instead of purchasing or leasing more in-house equipment that may be underutilized, customers simply pay for the cycles and compute or storage resources that they reserve for a specified period of time.

IBM provides a scalable, highly secure and extensible environment designed to handle peak workloads. This solution combines the best of IBM technologies and open standards with on demand computing, providing an infrastructure that enables customers to help meet their goals of increased scalability, flexibility, and processing capacity.

Building a world-class infrastructure to support the financial services industry

IBM has a comprehensive approach to deliver on demand computing for financial institutions. IBM's advanced clustering capabilities and services can help reduce the time, effort and expense required to analyze, simulate, and optimize a broad spectrum of risk management and compliance challenges.

The HPC infrastructure is based on the IBM @server™ Cluster 1350, an innovative Linux-based cluster that provides a robust, highly scalable compute facility centrally managed and controlled by IBM. The Cluster 1350 combines the power of IBM @server xSeries® Intel Xeon™ processor 32-bit and IBM @server AMD Opteron™ processor 32-bit/64-bit compute servers with IBM Cluster Systems Management (CSM) software, IBM TotalStorage® products, and leading third-party networking components to enable powerful, flexible solutions for high performance computing. IBM also offers clusters of IBM @server pSeries® POWER™ 64-bit compute servers running AIX 5L™ or Linux. IBM plans to offer other models, platforms and technologies over time.

Customers can access virtual dedicated Cluster 1350 cluster capacity based on the xSeries 335 2-way Intel Xeon processor-based server or IBM @server 325 2-way AMD Opteron processor-based server. The xSeries 345 or IBM @server 325 is used as a management node for centralized control of the system. Each virtual cluster has its own dedicated management node. Customers can also access virtual cluster capacity based on the pSeries 655 4-way server with a pSeries management node. Optional Fibre Channel or SCSI external disk storage capacity can be provided. Customers also have access to an IBM 3590, LTO, or DLT serially shared tape server for loading data and performing backups.

Customers have full control of the compute and storage resources within their assigned environment, and each virtual cluster has a dedicated cluster and management network. A robust networking infrastructure is designed to allow customer data and applications to be highly available and secure, and a virtual private network (VPN) service provides remote access to the facility. Customers pay for their assigned compute and storage capacity for the length of time that they use it.

Advanced systems management and Grid-enabling software

IBM Cluster Systems Management (CSM) for Linux software provides robust capabilities for resource monitoring, automated operations, remote hardware control, distributed command execution, configuration file management and parallel network installation. Access to the system is accomplished through a remote shell interface. Customers may use CSM directly to perform remote management tasks, install and use their own management tools, or request comprehensive management and monitoring services from IBM.

The IBM General Parallel File System (GPFS) for Linux and AIX® is a high performance scalable file system—ideal for the financial services industries that manage extremely large data files. GPFS can provide an enhanced layer of scalability, availability and performance. Optional dynamic job scheduling and workload management tools are also planned to be available to optimize cluster resources and help increase job throughput.

Customers supply the additional software required to run their workloads; this software may include in-house, open source, or IBM and third party applications and tools. Customers can use Grid software and services, acquired separately from IBM (such

as the 'IBM Grid Offering for Analytics Acceleration') or a third party, to develop new applications or run existing applications in a Grid environment. The open source Globus Toolkit™ is designed to provide the underlying elements for Grid security, communication, information infrastructure, resource management and portability. The IBM Grid Toolbox for Linux and AIX is an integrated set of tools and software based on the Globus Toolkit that facilitates the building of Grids and development of Grid applications. In addition to running production workloads, customers may run benchmarks or perform other types of analysis.

Services for delivering HPC on demand

IBM offers a variety of services to support IBM-supplied hardware, software, and networking components and infrastructure. Base services include the facility, provisioning and deployment of resources, security, virtual private network, monitoring, maintenance, and help desk support. Customers can benefit from a range of fee-based custom services, such as enhanced VPN bandwidth, software deployment and customization, advanced monitoring and management, storage management and backup, Grid-related services, and more.

Delivering dynamic HPC resources

IBM offers dynamic HPC resources on demand that are designed to match customer requirements for scalability and capacity. The on demand computing model is economical, enabling customers to enjoy flexible pricing and terms and helping them to realize lower operational and infrastructure fixed costs. This computing model also can help reduce infrastructure complexity, IT training and staffing and ongoing maintenance and upgrades. Customers can focus on their core business issues and be insulated from rapid information technology shifts.

IBM is extending its leadership in high performance computing to the on demand environment. By combining leading IBM @server technologies, advanced cluster and systems management capabilities, growth and innovation in Grid and autonomic computing, and demonstrated industry expertise, IBM is delivering solutions designed to enable financial services companies to conduct their analysis, simulation and optimization more efficiently and effectively in a highly competitive market.

Figure A. Possible technical components: Deep Computing Capacity on Demand for the Financial Services Industry

Component	Features/Benefits
HARDWARE	
IBM @server Cluster 1350	• Combines IBM @server servers, IBM TotalStorage storage and leading third-party networking components to create powerful, flexible solutions for HPC and commercial application environments • Uses xSeries 335 (1U) and xSeries 345 Intel Xeon 32-bit dual processor nodes • Uses BladeCenter Intel Xeon 32-bit dual processor nodes • Uses IBM @server 325 (1U) AMD Opteron 32-bit/64-bit dual processor nodes • Runs the Linux operating system standard • Runs the Windows operating system by exception • Uses CSM for central management
AMD Opteron Technology	• Runs 32-bit and 64-bit applications simultaneously • Easy migration to 64-bit computing capability with software investment protection • Extreme price/performance design—high performance at affordable pricing • Runs Linux and Windows
IBM pSeries Technology	• Features POWER 64-bit processors in an ultra-dense packaged server for high performance computing • Includes the AIX 5L operating system and a rich suite of cluster software
IBM TotalStorage	• FAStT Fibre Channel and SCSI disk options • 3590, LTO, and DLT tape server and cartridges options
SOFTWARE	
Operating Systems	• Linux and Windows provided by customer • AIX licenses provided by IBM
IBM Cluster Systems Management (CSM)	• Provides robust, powerful management from a central point of control • Simplifies administrative tasks and may reduce life-cycle costs • Offers a highly reliable infrastructure and event monitoring • Provides software installation and updates, remote hardware control, distributed command execution, configuration file management, and diagnostics
IBM General Parallel File System (GPFS) (optional)	• Provides shared access to files across multiple disk drives on multiple nodes • Provides a common file system abstraction for data shared among multiple nodes • Allows applications to easily access files using standard POSIX (Portable Operating System Interface for UNIX®) file system interfaces • Enables parallel applications to simultaneously access either the same or different files • Provides high availability through automatic recovery from node and disk failures
Job scheduling (optional) (planned availability)	• Provides dynamic job scheduling and workload management • Designed to optimize cluster resources and increase job throughput • Works with CSM to facilitate management of cluster resources
SERVICES	
IBM Global Services	• Support Deep Computing Capacity on Demand with a full range of services for solution customization • Provide Grid enablement services, including Grid Innovation Workshops, Grid Pilot Implementation Services, and Grid Rollout Implementation Services

For more information

To learn more about IBM Deep Computing Capacity on Demand for the Financial Services industry, contact your IBM representative.

AMERICAS

Alan McCarter

e-mail: armccar@us.ibm.com

ASIA PACIFIC

Sinisa Nikolic

e-mail: 9sinisa@my.ibm.com

EUROPE/MIDDLE EAST/AFRICA

Ian Green

e-mail: ian_green@uk.ibm.com

IBM Corporation
Systems Group
Route 100
Somers, NY 10589

Produced in the United States of America
January 2004

GM13-0517-00