

## IBM Grid Offering for Analytics Acceleration: Financial Markets



### Highlights

- *Helps enhance competitiveness and agility in the financial trading market*
- *Reduces statistical margin of error, speeds trade decisions and increases number of scenarios and parameter space*
- *Offers affordable, effective IT for a sophisticated trade portfolio, and tolerance for IT asset failure*

### The challenges facing financial analysis

In recent years, the science of financial engineering has permeated the banking and finance industries, giving birth to thousands of new products and techniques for financial analysis.

Pricing financial risks accurately is a core competency in a very competitive marketplace. Along with sophistication and accuracy, speed is a primary success factor. Offering the best product at the right time determines the winners in the race for the best margins.

Therefore, applying pricing and simulation analytics faster than the competition—in a highly dynamic environment where new products appear every day and activity shifts rapidly between asset classes—is a major challenge. It demands that the power, reactivity and flexibility of the technology infrastructure be combined in the most efficient fashion. These IT environments must also be able to embrace new algorithms and techniques while being adaptable, extendable and scalable.

### Introducing Grid Computing

Grid Computing is a services-oriented architectural approach that uses open standards to enable distributed computing over the Internet or a private network. This approach helps IT departments aggregate disparate elements such as compute resources, data storage and filing systems to create a single, unified system.



## Grid Computing

At its core, Grid Computing enables devices—regardless of their operating characteristics—to be virtually shared, managed and accessed across an enterprise, industry or workgroup. This resource virtualization provides the necessary access, data and processing power to rapidly solve complex business problems, conduct compute-intensive research and data analysis and engage in real-time business On Demand. Grid Computing helps to promote the efficient utilization of technology resources and fosters the creation of cost-effective, resilient IT infrastructures that are adaptable to change.

### **Grid Computing and analytics acceleration in financial markets**

Today's financial trading environment requires fast, accurate decision making to help optimize the investment profile and manage the associated risk. The shortest time taken to make optimal decisions is key to gaining a competitive edge in the investment market. Grid Computing helps financial institutions accelerate their trading analytics operations and increase computational throughput.

Businesses can use untapped compute and processing resources to considerably accelerate analytical processes—reaching end results far more rapidly than within conventional computing environments. One clear advantage of a distributed environment for financial analysis operations is the ability to run several trade scenarios in parallel to generate an optimal solution as quickly as possible, without sacrificing increased accuracy and performance.

Depending on the customer's analytics applications, flexibility requirements and existing infrastructure, two other techniques can be used in combination with Grid Computing to facilitate analytics acceleration: application parallelization and hardware acceleration. Application parallelization divides an application into independent tasks that can then be executed concurrently, thereby reducing the time to completion. Hardware acceleration applies additional processing power by establishing or growing a cluster solution, or by using faster processors.

### **IBM Grid Offering for Analytics Acceleration**

IBM has a comprehensive approach to help determine the most appropriate combination of technologies for analytics acceleration in financial markets. Award-winning WebSphere® Web application servers, advanced computer cluster offerings, powerful database technologies, application integration approaches and knowledge management techniques from IBM can help reduce the time and effort required to capture, compile and analyze research data and can accelerate product development. IBM Grid Offerings include a Grid Innovation Workshop, a pilot implementation and a complete portfolio of IBM Global Services offerings.

### **IBM product foundation: supported platforms for Grid Computing**

IBM offers several products to build a Grid Computing platform for analytics acceleration, including IBM @server pSeries™ servers, and the IBM @server Cluster 1600 and IBM @server Cluster 1350 platforms.

Innovative pSeries servers are designed to lower costs, improve efficiency and speed the transformation to a Grid-enabled infrastructure. The Cluster 1600 brings together cluster-ready pSeries servers, cluster management software and a choice of IBM or industry-standard cluster interconnect technologies. The Cluster 1350 combines the power of rack-optimized IBM **@server** xSeries™ servers with IBM Cluster Systems Management for Linux®, optional IBM storage products and leading third-party networking components to create powerful, flexible solutions for high-performance computing and commercial application environments.

### **IBM product foundation: Grid-enabling software**

The open source Globus Toolkit™, developed through The Globus Project™, includes a full range of Grid-based commands and programming tools to quickly develop Grid applications or to run existing applications in a Grid environment. It

provides the underlying elements for Grid security, communication, information infrastructure, resource management and portability.

The IBM Grid Toolbox is based on the Globus Toolkit and includes enhanced documentation, installation tools for AIX® and Linux environments and optional IBM service and support offerings. The Toolbox can be deployed on various IBM **@server** products, including pSeries and xSeries servers as well as IBM Cluster 1350 and Cluster 1600 platforms.

See Figure A, which outlines available technologies and IBM software that enable Grid Computing for analytics acceleration.

### **ISV Grid technologies**


IBM has established strong relationships with Grid middleware independent software vendors (ISVs). The technologies provided by these ISVs offer innovative solutions to some of today's business problems. IBM has

assessed these technologies in detail, helping to ensure that they adopt appropriate standards and provide interoperability, scalability and resilience.

### **Parallelization software**

Application parallelization can facilitate Grid Computing when used for analytics acceleration. IBM uses the Message Passing Interface (MPI) or Parallel Virtual Machine (PVM) to help parallelize your applications.

Clearly, parallelizing the application is just the first step—running the code on a distributed set of virtualized resources is required to significantly increase the computational throughput. IBM's relationships with Platform Computing, DataSynapse, Avaki and United Devices allow you to use a single product for application parallelization and distributed execution. These four companies offer products that provide job- and task-level scheduling for low-overhead job submission as well as for interactive





## Grid Computing

computations. IBM also offers LoadLeveler® with Globus Toolkit 2.0 support, a high-performance computing cluster job scheduler for xSeries and pSeries servers, which provides the advantages of checkpoint restart and backfill scheduling.

### **IBM services to support Grid technologies**

Grid Computing can be adapted to serve unique business environments, and IBM Global Services (IGS) supports all elements of this architecture—servers, operating systems, middleware and networks—with a full range of Grid-related services. For example, IGS can provide a business-case analysis, cost justification and a strategic perspective for leveraging Grid Computing across the enterprise. For other customers, IGS may review potential applications, select the appropriate technologies and help enable the selected applications to use Grid Computing effectively.

IGS also can design and deploy Grid infrastructures that provide the required performance, scalability, resiliency and reduced operational costs. Or, if you need an end-to-end Grid solution, IGS can integrate all of these capabilities into a coordinated approach that delivers a comprehensive Grid implementation.

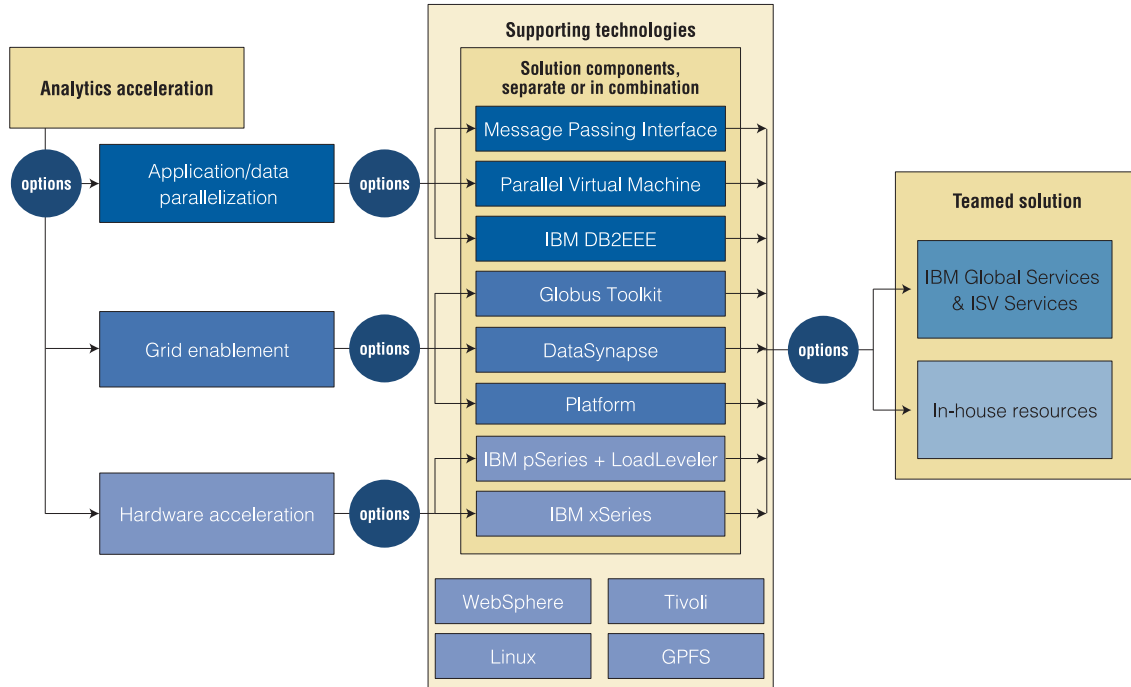
Furthermore, IBM's array of strategic relationships with Grid technology developers allows you to review and select the appropriate technologies. By offering these options and capabilities, IGS can provide you with a low-risk, efficient and effective Grid solution.

### **IBM as a leader in Grid Computing**


IBM is committed to open standards and is working with The Globus Project open source development community, the Global Grid Forum (GGF) and the Interoperable Informatics Infrastructure Consortium (I3C) to promote the adoption of open standards and accelerate the availability of Grid technology.

IBM is a strong supporter of The Globus Project, a multi-institutional research and development effort to address the technical and business challenges of Grid Computing. Founded by a team of technicians and researchers, The Globus Project has defined an open source Grid reference architecture and a set of tools to assist Grid deployment.

IBM Design Centers use the latest Grid technologies, including the Open Grid Services Architecture (OGSA), which merges the open protocols used for Grid Computing with the protocols used for Web services. IBM Design Centers offer customers access to the latest software from leading Grid software companies, such as Data Synapse and Platform Computing, and the latest open source Grid technologies from The Globus Project.



**Figure A. Analytics acceleration decision landscape**






## Grid Computing

**Figure B. Engagement points for Grid enablement**

<b>Grid Innovation Workshop for Financial Markets</b>	<b>Grid Pilot Implementation Services</b>	<b>Grid Rollout Implementation Services</b>
<p>IBM can help you decide which option is right for you. When appropriate, IBM Global Services Grid Innovation Workshop can help to develop the vision, plan and execution model for your Grid project.</p>	<p>IBM can help define a pilot implementation project that will help you to validate the benefits of Grid Computing. Piloting services include:</p> <ul style="list-style-type: none"> <li>• Introduction, education, installation and support of the Globus Toolkit</li> <li>• Selection of Grid middleware partners relevant to your needs</li> <li>• Training on Grid concepts and implementation techniques</li> <li>• Creation of a test Grid</li> <li>• Access to public data sources</li> <li>• Projection of the benefits of a full Grid implementation</li> <li>• Timely and significant education and analysis of the data Grid technology</li> </ul>	<p>IBM Global Services can assist with your Grid implementation project by providing an array of implementation services such as:</p> <ul style="list-style-type: none"> <li>• Project management support</li> <li>• Business-process consulting services and training</li> <li>• Help desk</li> <li>• Data integration services through federation</li> <li>• Portal development</li> <li>• Linux implementation services</li> <li>• High-performance clusters</li> <li>• Engagement services</li> <li>• IBM Security Services</li> <li>• Systems integration services</li> </ul>

**Figure C. Possible technical components: IBM Grid Offering for analytics acceleration in financial markets**

Feature	Benefits
<b>PARALLELIZATION</b>	
<b>Message Passing Interface (MPI)</b>	<ul style="list-style-type: none"> <li>• Uses open standards and helps ensure portability over many platforms</li> <li>• Provides point-to-point communication, collective communication, virtual process topologies and support for writing heterogeneous parallel applications</li> <li>• Can be attained as either a shareware or vendor-supplied implementation</li> <li>• Can provide significant flexibility in performance tuning and programming or in a message-passing environment</li> </ul>
<b>Parallel Virtual Machine (PVM)</b>	<ul style="list-style-type: none"> <li>• Uses open source software</li> <li>• Can be used with heterogeneous systems (running the UNIX® or Microsoft® Windows® operating systems)</li> <li>• Permits systems to be coupled together across a network and used as a single, large parallel computer</li> </ul>
<b>IBM DB2® Universal Database Enterprise—Extended Edition (DB2EEE)</b>	<ul style="list-style-type: none"> <li>• Provides shared-nothing architecture in which each machine in a cluster has exclusive access to its own disks and memory and does not compete with other machines for resources</li> <li>• Facilitates complex parallel processing operations on platforms such as AIX, HP-UX, Microsoft Windows NT®, Linux and the Solaris™ Operating Environment</li> </ul>
<b>GRID</b>	
<b>GLOBUS Toolkit/IBM Grid Toolbox</b>	<ul style="list-style-type: none"> <li>• Facilitates quick development of Grid applications</li> <li>• Can be used with various IBM @server products</li> </ul>
<b>General Parallel File System (GPFS)</b>	<ul style="list-style-type: none"> <li>• Provides shared access to files across multiple disk drives on multiple nodes</li> <li>• Provides a common file system abstraction for data shared among all the nodes in a Grid</li> <li>• Allows applications to easily access files using standard UNIX file system interfaces</li> <li>• Allows parallel applications to simultaneously access either the same or different files</li> </ul>
<b>WebSphere platform</b>	<ul style="list-style-type: none"> <li>• Provides infrastructure software for dynamic e-business. Founded on open standards such as J2EE and Web services, WebSphere has evolved from a Web application server to a full set of related products and offerings with a common base. IBM expects WebSphere to be the hosting environment for customers wishing to exploit the Open Grid Services Architecture</li> <li>• Delivers value in three areas: <ul style="list-style-type: none"> <li>• <b>Reach and user experience.</b> Enhances customer loyalty by delivering a rich user experience through offerings such as WebSphere Portal</li> <li>• <b>Business integration.</b> Helps slash process cycle time and costs, and thus can help improve return on investment (ROI), through offerings such as WebSphere MQ</li> <li>• <b>Foundation and tools.</b> Provides a reliable, scalable and flexible infrastructure for dynamic e-business through offerings such as WebSphere Application Server and WebSphere Studio</li> </ul> </li> </ul>






## Grid Computing

Feature	Benefits
<b>GRID</b>	
<b>Tivoli®</b>	<ul style="list-style-type: none"> <li>• Provides highly scalable, cross-platform systems management</li> <li>• Enables efficient management of e-business infrastructures from a single location</li> <li>• Can help maximize ROI and speed deployment</li> <li>• Can simplify management of separate devices without sacrificing productivity, security or performance</li> </ul>
<b>ISV Grid technologies</b>	• Offer innovative solutions to today's business problems
<b>Linux</b>	• Provides an open source code base for Grid development efforts
<b>IBM LoadLeveler</b>	<ul style="list-style-type: none"> <li>• Provides a high-performance computing cluster job scheduler for xSeries and pSeries servers</li> <li>• Enables checkpoint restart and backfill scheduling</li> </ul>
<b>HARDWARE</b>	
<b>IBM @server pSeries</b>	<ul style="list-style-type: none"> <li>• Can help to lower costs, improve efficiency and speed transformation to a Grid-enabled infrastructure</li> <li>• <b>pSeries 690:</b> Features two 1.1 or 1.3GHz POWER4™ microprocessors, the first “server-on-a-chip” design for high-end UNIX servers, with up to 32 processors per server; delivers true logical partitioning; can be divided into as many as 16 virtual servers, each with its own set of system resources such as processors, memory and I/O</li> <li>• <b>pSeries 670:</b> Features one 1.1GHz POWER4 processor, with up to 16 processors per server and the same logical partitioning capabilities as the pSeries 690</li> <li>• <b>pSeries 630-6C4:</b> Features up to four 1GHz POWER4 processors; offers the lowest entry price for a pSeries four-way cluster building block</li> <li>• <b>pSeries 660:</b> Features a range of three types of powerful midrange, rack-mounted servers; offers optimal flexibility and outstanding power in a small footprint</li> <li>• <b>pSeries 655:</b> Incorporates up to 16 servers, with up to 128 POWER4 processors, into an ultra-dense packaged server</li> </ul>
<b>IBM @server Cluster 1600</b>	<ul style="list-style-type: none"> <li>• Provides a highly scalable AIX cluster and Grid Computing solution for large-scale computational modeling, multi-terabyte databases and cost-effective data center, server and workload consolidation</li> <li>• Brings together cluster-ready pSeries servers, cluster management software and a choice of IBM or industry-standard cluster interconnect technologies</li> <li>• Extends the benefits of IBM RS/6000® SP™ technology to additional hardware building blocks</li> <li>• Can be built using the latest, most powerful pSeries servers, including pSeries 690, 670 and 630 servers, to deliver high-performance computing solutions and reduce complexity and total cost of ownership for managing multiple systems</li> <li>• Provides a choice of Parallel System Support Programs (PSSP) or Clusters System Management (CSM) for single-point-of-management control</li> <li>• Provides a choice of industry-standard (10/100Mbps or 1Gbps Ethernet), SP Switch or SP Switch2 interconnect technologies</li> <li>• Includes the AIX 5L operating system and a complete cluster software suite for high-end technical computing</li> <li>• Includes High Availability Cluster Multiprocessing (HACMP) for AIX, which is designed to keep business-critical applications and systems operational 24x7. HACMP can automatically detect system and application failures; provide the capability to recover system hardware, applications, data and users while keeping recovery time to an absolute minimum; and enable hardware and software upgrades without interrupting cluster operations</li> </ul>



Feature	Benefits
<b>HARDWARE</b>	
<b>IBM @server Cluster 1350</b>	<ul style="list-style-type: none"> <li>• Runs on the Linux operating system</li> <li>• Can operate in a standalone environment for dedicated application workloads or in a Grid Computing environment for high-throughput resources and high performance</li> <li>• Combines rack-optimized IBM @server xSeries servers, CSM for Linux, optional IBM storage products and leading third-party networking components to create powerful, flexible solutions for high-performance computing and commercial application environments</li> <li>• Can be used in a wide range of configurations, including single-rack systems of up to 32 cluster nodes and large-scale systems of up to 512 cluster nodes</li> <li>• Uses xSeries 335 or xSeries 345 as the cluster nodes: one- or two-way x335 servers (1U), with Intel® Xeon™ processors running at either 2.0 or 2.4GHz; may optionally include up to 32 one- or two-way x345 storage servers (2U), with Intel Xeon processors running at 2.0 or 2.4GHz. Each cluster requires a two-way x345 management node</li> <li>• Can be managed by CSM for Linux Version 1.2 or 1.3, which provides resource monitoring, automated operations, remote hardware control and command execution, configuration file management and parallel network installation; CSM for Linux helps ease administration and may reduce life-cycle costs by allowing management of an entire Cluster 1350 system from a single management node</li> </ul>
<b>SERVICES</b>	
<b>IBM Global Services</b>	<ul style="list-style-type: none"> <li>• Supports all elements of Grid Computing with a full range of services</li> <li>• Allows customers to review and select the appropriate technologies</li> </ul>



## For more information

To learn more about Grid Computing for analytics acceleration in financial markets, contact your IBM representative or IBM Business Partner, or visit [www.ibm.com/grid](http://www.ibm.com/grid).

### AMERICAS

Basu Hurkadli  
e-mail: [bnhurka@us.ibm.com](mailto:bnhurka@us.ibm.com)

### ASIA PACIFIC

Takayuki Takano  
e-mail: [ttakano@jp.ibm.com](mailto:ttakano@jp.ibm.com)

### EUROPE/MIDDLE EAST/AFRICA

e-mail: [Philippe Bricard](mailto:Philippe.Bricard@fr.ibm.com)  
[bricard@fr.ibm.com](mailto:bricard@fr.ibm.com)



© Copyright IBM Corporation 2003

IBM Corporation  
Integrated Marketing Communications  
Server Group  
Route 100  
Somers, NY 10589

Produced in the United States of America  
01-03  
All Rights Reserved

IBM, the IBM logo, the e-business logo, AIX, DB2, LoadLeveler, POWER4, pSeries, RS/6000, SP, Tivoli, WebSphere and xSeries are trademarks of International Business Corporation in the United States, other countries, or both.

Globus Toolkit and The Globus Project are trademarks of the University of Chicago.

Intel and Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Solaris is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

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, our warranty terms apply.

All information in these materials is subject to change without notice. ALL INFORMATION IS PROVIDED ON AN "AS IS" BASIS, WITHOUT ANY WARRANTY OF ANY KIND.