



Grid Computing in Financial Markets: Moving Beyond Compute-Intensive Applications

November 2003
ES-01:001

Executive Summary

We acquire technology. Be that as a function of our acquisitive nature, the development of new applications, our need for five 9s reliability, or for equipping new staff. We buy desktops, servers, mid-range and mainframe computers. However do we really fully utilize them?

The answer to that question for most firms is of course, no. How much computing power is really needed to use Word or PowerPoint? How much computing power is needed to enter a trade, update a customer file, or send a single FIX message? Not much, but this does not stop us from acquiring hardware as we build more and more solutions. There have been many studies that have analyzed corporate technology utilization and concluded that most firms are lucky if their corporate server utilization reaches 15%. If you add on desktop utilization, it would be surprising if any large financial institution's average utilization rate hit more than 5%. However, just because a company only utilizes 5% of their compute power doesn't mean they don't need more capacity. If we can't process peak transaction loads during market open, close, or during volatile times, someone will be jobless.

But how do we utilize all of this excess capacity without being short cycles during the 10 billion-share trading day, or when we analyze enterprise risk on a global basis? Our hardware is heterogeneous. We run multiple operating systems and our applications are written and optimized for specific hardware and software platforms.

While our hardware and operating environment is heterogeneous, our application suite is becoming more component-based. Many of the newer programming languages and data protocols such as Java, XML and Web services are focused on making objects more robust. The languages are also concentrated on turning traditional applications that are centrally hosted into services that can be deployed virtually anywhere on the network. However there is a huge gap between being *capable* of deploying and *actually* deploying the service anywhere on the network.

To deploy a Web service, there needs to be a comfort level, a guarantee that the service will deploy to an active location, it has the appropriate data, it will execute properly, it won't be

compromised and the result will be integrated into the appropriate location, or, if part of a process, be connected to the appropriate back-end. But how do you control a heterogeneous technology and operating environment when you don't know which machine you will deploy on or even the operating environment that will be processing your calculations? This is the challenge and the promise of grid computing.

The promise of grid computing is more than just taking large and as my undergraduate statistics professor liked to say "ugly" computationally intensive problems, sub-divide them in to smaller tasks, parse them out to a bank of idle PCs and retrieve the answer back relatively quickly at a relatively low cost.

The promise of grid computing guarantees the ability to virtualize, or make transparent, the network, compute and data infrastructure from the application environment that can operate wherever and whenever it makes sense. Grid computing allows the possibility to leverage the proper hardware, in the proper location, for the appropriate process, transparently. It will enable applications to automatically split large computational problems into smaller pieces and run them on idle servers and desktops, or break up an application to smaller components and run as many instances of those components as needed. In context, the promise of grid in a trading environment would be to automatically provision and instantiate more order management servers during a busy market open or close and re-provision them to risk management or back-office functions when the market calms down.

Why Is Grid Hot? Grid Technology Business Drivers

Grid is hot because it's the right technology for the time. Cost, complexity and opportunity are driving this market. Cost, now more than ever has been a tremendous issue for financial markets firms. The challenging revenues and market conditions have made financial markets firms look closely at expenditures, most notably technology expenditures, which prior to 2000 were escalating at rates greater than 15% per year. In firms' attempts to clamp down on their rampant technology expenditures and the post-Internet bubble revenue squeeze, they began to look at the cost of technology as well as its utilization. Since many applications are directly linked to their hardware and storage infrastructure and firms configure systems for peak utilization, general technology utilization rates for many platforms are very low. However as utilization increases and more headroom is needed, the cost of scaling applications on large SMP computers can escalate, especially since the cost of large SMP boxes increases non-proportionally to the number of CPUs in the box.

This has pushed many firms toward using low cost Intel-based clusters which can be scaled incrementally whereas a large SMP box utilizing all of its CPU slots requires the purchase of a new larger capacity SMP box (plus another for disaster continuity services) instead of incrementally increasing the size of a cluster through the purchase of a few more inexpensive nodes.

Grid technologies promises to link these disparate technologies and clusters together so they can be easily repurposed and better utilized, thus reducing the overall purchase of hardware, adding utilization and increasing scalability of existing applications.

The increasing complexity of firms' technology infrastructures also is an argument for the implementation of grid technologies as the increasing heterogeneous nature of firms'

architectures makes it more difficult to manage. Grid technology promises to make technology management more virtual and have a single or more limited selection of management tools that can manage a wider diversity of technologies than dedicated management tools.

The other major driver forcing firms down the grid path is the opportunity to perform tasks that were impractical or impossible to do without a massive compute infrastructure and the opportunity to better leverage component-based architectures (Web services) to extend technology and services to clients as well as to better utilize common functions and services within an organization.

Financial markets are information based. Much of this information is created, stored and analyzed to better understand the markets, internal risk and potential opportunities. To analyze this information, firms need an increasing array of technology. Unfortunately, firms these days do not have an increasing amount of funds to acquire technology to tackle these queries. Grid enables firms to better utilize the resources that they currently have and expedite the analytics they perform on a regular as well as an ad-hoc basis. The ability to provide these analyses in a faster timeframe by using grid computing allows the opportunity to take advantage of this information before competitors.

Firms are also looking to leverage their own technologies better and extend these technologies to their clients via Web services. The problem with Web services has been the difficulty of both linking together services to create complex processes as well as managing the deployment of these services over a heterogeneous technology infrastructure. Grid technology vendors are working to accomplish both of these feats to ensure that complex processes can be deployed and managed with an acceptable quality of service.

The Tabb Group Report on Grid Computing in Financial Services

This report discusses the changing grid solutions market. It investigates the drivers, the taxonomy and vendors that are developing cutting-edge solutions in this rapidly changing environment. The report analyzes new grid standards and their impact on financial markets, the existing grid computing vendors and how firms are adopting and adapting to this new grid vision. This report provides an in depth analysis of grid computing adoption, market size projections and estimates the project growth levels of grid computing in financial markets. A table of contents can be found below.

A copy of this report is available at www.tabbgroup.com/research.

Table of Contents

Vision	1
Introduction	4
What Is Grid Computing?	5
Grid, Clusters and Parallel Computing	5
Why Is Grid Hot? Grid Technology Business Drivers	6
Grid Taxonomy	7
Scale	7
Function	8
Batch Grids	11
Grid, Utility and Autonomic Computing	11
Grid Components	12
Grid Standards	13
Web Services, BPEL and Peer to Peer: The Next Frontier	14
BPEL & Service Grids	15
Peer to Peer	16
Grid in Financial Markets	17
Computational Grids	17
Data Grids	18
Services Grids and Technology Management	19
One Firm's Frustration—A Case Study	19
Challenges	20
Moving Away from Cost as a Value Proposition	21
Market Sizing	22
Financial Markets Grid Vendors	23
Platform Computing	24
Data Synapse	25
United Devices	26
Data Grid Providers	27
Avaki	27
GemStone	27
MetaMatrix	28
Oracle	28
Integrasoft	29
Service Grids	29
Enigmatec Corporation	29
The Globus Toolkit	31
The Three Kings: IBM, HP and Sun	31
Conclusions	33

