



Platform[™]
The Power of Sharing

ROI of Utility Computing





"Our enterprise grid is enabling our IT users to be more creative in how they share computing resources to do their development and testing work, thus helping us achieve our expensive efficiency targets."

John van Uden, SVP, Citi

As the executives of every financial institution know very well, demand for compute power continues to grow. Not only are risk calculations, Monte Carlo simulations, and pricing models getting more complex, but each financial institution wants to reduce the amount of time it takes to make the calculations.

Since many firms pursue similar trading strategies, speed becomes a differentiator. Executives worry that excessive latency in the execution of trading applications leads to a competitive disadvantage by skewing market pricing or resulting in trade rejections.

These complex trading applications include functionality such as trade order capture and execution, pricing engines, and data feeds from sources such as Reuters. The more inputs and components the greater the potential penalty on speed. Over the course of time, a few milliseconds of latency add up. Those that are empowered to see new opportunities first and capitalize on them before the competition will make the biggest profits.

One computer, two computer...

For many institutions in the financial services (FS) industry the preferred strategy to support complex distributed applications has been to buy speed by spending more money:

- buy more hardware
- hire more people to manage and support the growing clusters
- silo each new application to guarantee computing resources and protect SLAs

Cluster after cluster is built to power mission-critical applications. Unfortunately, a company that silos its applications in separate lines of businesses inherently reduces its competitive edge. The logic is simple: each new application silo has to be equipped to handle the theoretical peak demand of the application. Of course peak demand may only be reached at certain times of the day or week, if ever. The result is a low average utilization rate and lots of wasted capacity.

Three computer, four...

How much wasted capacity? A recent study of six corporate data centers revealed that most of their servers were using less than a quarter of their available processing power. The replication of several independent clusters, all using similar hardware, running similar software, and employing similar kinds of workers, imposes severe penalties on a company's bottom line. It leads to the overburdening of IT assets, dampening the productivity gains that can spring from computer automation.

Is there something I can do about it?

The simple answer is move to a utility computing model where many departments and many applications share a pool of compute resources. In order to do so, you need to navigate the political landscape in your organization in order to get buy in. Our recommendation is to start with a handful of applications on a shared grid and once you have proved the concept, it becomes easier to convince the organization of the value.

What makes this scenario possible is Platform Symphony, a key product of Platform Accelerate suite, one of two suites in Platform's High Performance Computing (HPC) management software enterprise offering. Platform Symphony manages and regulates the traffic on the grid, making sure that application processes are distributed efficiently across available clusters and individual nodes. The software ensures that each application has access to enough resources for maximum performance.

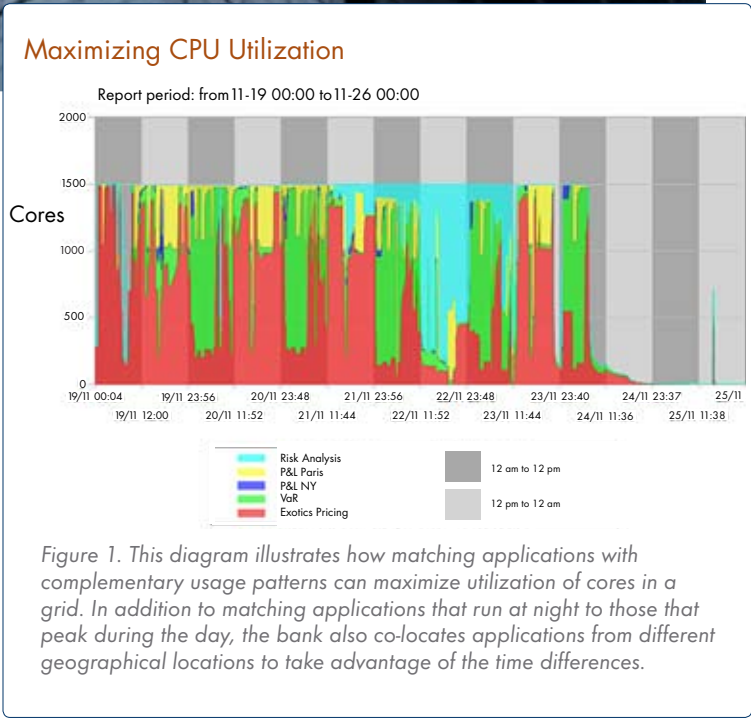


Figure 1. This diagram illustrates how matching applications with complementary usage patterns can maximize utilization of cores in a grid. In addition to matching applications that run at night to those that peak during the day, the bank also co-locates applications from different geographical locations to take advantage of the time differences.

Cost comparison over one year – silo grids vs. utility grid				
	Data Center	Silo Grids	Utility Grid	Area of Saving
	# of CPUs	4,000	2,500	Server Consolidation – resource sharing
CAPEX	Hardware + Chassis	\$20,000,000	\$12,500,000	Move to more commodity hardware
	OS support	\$400,000	\$250,000	Reduced OS Subscriptions
	Storage	\$86,400	\$86,400	No effect
OPEX	Data Center (power, cooling)	\$1,822,080	\$1,138,800	Utility Grid = Green Computing
	Real Estate	\$600,000	\$600,000	No effect
	System Administrator	\$2,500,000	\$1,750,000	More CPUs managed per Admin
	Application Dev / Support team	\$4,000,000	\$4,000,000	No effect
	Total	\$29,408,480	\$20,325,200	

Figure 2. This chart shows a simplified calculation of the potential savings and costs deferrals associated with a utility grid running two analytical applications. This example shows a \$9 million US savings. With each additional application added to the grid the savings grow.

Platform Symphony keeps the entire system working as speedily and efficiently as possible, scheduling analytic calculations to run on nodes with the required hardware and software components, and re-routing processes around failed nodes. The beauty of the utility grid is that it is essentially failsafe; no one component is essential to the system.

Utility computing in the real world

The benefits realized by two Platform customers in the financial services industry that have taken full advantage of HPC utility computing are described in the following paragraphs. The first customer built a utility grid that combines 15 major intra- and end-of-day pricing and risk applications. In the process, they were able to drive CPU utilization from an average of 20% to over 80%. The new global utility platform in place, powered by Platform Symphony, provides on-demand processing power while helping the customer defer significant hardware costs.

The second customer used Platform Symphony to build a utility grid that runs seven major trading and risk systems. They calculate that this has saved the organization several million dollars. One significant area of savings is the cost of powering, cooling and maintaining the nodes in the grid. According to an article in *American Banker*, this customer was able to reduce its cost per CPU to around half of the industry standard.

Both organizations have been able to navigate the pitfalls of inter-departmental politics and territorialism in order to roll-out their utility grids.

Performance at any scale with low latency

Platform Symphony is the fastest, most fault-tolerant and scalable HPC service-oriented middleware solution, able to perform analytical processing tasks in under two milliseconds. A component of Platform Symphony, Platform Enterprise Grid Orchestrator (EGO) technology manages a shared pool of nodes to meet the demand of multiple applications in a utility model – enabling enterprise-class resource sharing, availability, security, and scale.

Proven scalability, utilization, performance	
	Symphony at IBM DCCoD
Scalability 1,000 concurrent clients, 100 applications	20,000+ CPU's simulated on 1,000 physical CPUs in one cluster
CPU Utilization 1-100 clients, 1 sec task, 1KB message, 2,000 CPU	98%
Single Task Latency	< 1 ms

Figure 3. Shows test results demonstrating the scalability, utilization and task throughput with Platform Symphony.

Technology enabling virtualized, shared resource pools

Platform technology also supports enterprise class resource sharing. Platform EGO is a resources broker that enables delivery of a shared, virtualized pool of IT resources, to multiple application types based on business policies. It is an application- and infrastructure-aware technology that allows Platform HPC management software to decouple resource demand from resource supply. This solves the datacenter business pain of under- or over-provisioning by translating unpredictable infinite demand to finite compute resources.

Cluster lifecycle management reduces costs

The value of Platform Manage, the second of two suites comprising Platform's HPC management software offering, lies in rapid, consistent, and automated cluster lifecycle management in a HPC datacenter. The two suites share a single centralized interface. Platform Manager, a key component of this suite, is designed to allow companies to follow a strategy of building heterogeneous, multi-vendor, HPC datacenters out of low cost commodity hardware.

Cluster management software simplifies cost-control strategies such as building clusters by sourcing affordable, commodity hardware. Whether deploying a new cluster from the bare metal up, or adding capacity to an existing cluster, this strategy has some built in entropy. The cluster can end up with a heterogeneous mix of hardware, operating system versions, and application software from a number of vendors.

Efficiently monitoring and managing such a heterogeneous cluster, the nodes, the application and operating system licenses, and the priority of jobs being queued to the cluster, adds great value to the process. Platform Manager plays a key role in these aspects of improving the performance of heterogeneous clusters, allowing companies to achieve significant savings by reducing the time spent on deploying, configuring, managing and operating clusters. Various other Platform Manage suite products also supply reporting, tracking and analytics tools to aid in understanding and improving cluster performance and planning for future use patterns.

The Green Grid

As a member of The Green Grid™, Platform Computing is prominent among software companies meeting the challenges of aligning environmental responsibility with business success. Platform Symphony helps make grids greener by reducing power consumption for running and cooling the clusters, as well as reducing the physical footprint they occupy. This is mostly achieved by driving up utilization rates as high as 98% so that you can do more work with fewer machines.

Leverage IT to drive innovation and performance

Platform offers a complete low-latency solution for building a utility grid. That it helps FS institutions make money through better, faster trades; save money by better cluster monitoring and management; and facilitate greener computing. Given recent economic developments, it is more important than ever for FS firms to leverage IT to help make faster, more accurate business decisions. This new found agility will help drive the performance and innovation the industry needs to evolve, while providing real, tangible benefits that will increase the ROI of your IT infrastructure and contribute to a greener IT strategy. All of these factors work together to help financial institutions grow their business while managing risk, in a cost effective, utility computing model.

The Platform Promise

Platform's HPC management software reduces the total cost of ownership (TCO) by fully utilizing existing hardware and facilitates the easy deployment and use of economical, heterogeneous, commodity hardware-based systems. Benefits include:

- defer hardware additions
- eliminate cost of maintaining multiple separate clusters
- more efficient usage of available application licenses
- reduced power consumption for heating and cooling
- efficiently redeploy human resources
- eliminate in-house software development
- reduce re-qualification of software to keep up with upgrades
- mitigate risk of knowledge loss through retirement of application developers

Find out more at www.platform.com.



Platform Customers

Fortis
Citi
Prudential
Lehman Brothers
RBC Capital Markets
UBS
Deutsche Bank
Mitsubishi UFJ Group
MassMutual Financial Group
Wells Fargo
Millenium Partners, L.P.
Prediction Company
KBC
Mizuho Financial
BNP

Trading and Risk Systems Partners

Calypso
Sungard
OpenLink
Misys
Numerix
Application Networks
Sophis
Codefarm

Insurance Partners

Axis GGY Inc.
Towers Perrin Tillinghast

Hardware & Software Vendor Partners

IBM
Intel
HP
Sun
Dell
Cisco
TopSpin
Oracle
Gemstone
GigaSpaces
SAS
Microsoft
Tivoli

Platform HPC Management Software

Platform offers a complete HPC management software solution that supports a single, consistent HPC environment. Its two suites, Platform Accelerate and Platform Manage, provide the tools you need to:

- manage the lifecycle of multiple clusters
- accelerate performance of compute or data intensive applications
- monitor, report on and analyze cluster performance
- manage license usage
- share resources across the enterprise

Platform Accelerate

Platform LSF
Platform LSF MultiCluster
Platform LSF License Scheduler
Platform LSF Session Scheduler
Platform Process Manager
Platform EGO
Platform MPI
EnginFrame

Platform Manage

Platform Manager
Platform Analytics
Platform RTM
Platform VM Orchestrator

Platform™

Platform Computing is the leader in grid and cloud computing software that dynamically connects IT resources to workload demand according to business policies. Over 2,000 of the world's largest organizations rely on our solutions to improve IT productivity and reduce data center costs. Platform has strategic relationships with Cray, Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat®, and SAS®. Building on 16 years of market leadership, Platform continues to help data centers be more efficient, responsive and dynamic. Visit www.platform.com.

World Headquarters
Platform Computing Inc.
3760 14th Avenue
Markham, Ontario
Canada L3R 3T7
Tel: +1 905 948 8448
Fax: +1 905 948 9975
Toll-free tel: 1 877 528 3676
info@platform.com

Sales - Headquarters
Toll-free tel: 1 877 710 4477
Tel: +1 905 948 8448

North America
New York: +1 646 290 5070
San Jose: +1 408 392 4900
Detroit: +1 248 359 7820

Europe
Basingstoke: +44 (0) 1256 883756
London: +44 (0) 20 7977 1480
Paris: +33 (0) 1 41 10 09 20
Düsseldorf: +49 2102 61039 0
Munich: +49 89 517397 52
info-europe@platform.com

Asia-Pacific
Beijing: +86 10 82276000
Xi'an: +86 029 87607400
asia@platform.com
Tokyo: +81 (0)3-6302-2901
info-japan@platform.com
Singapore: +65 6307 6590
wliaw@platform.com

Copyright © 2009 Platform Computing Corporation. The symbols ® and ™ designate trademarks of Platform Computing Corporation or identified third parties. All other logos and product names are the trademarks of their respective owners, errors and omissions excepted. Printed in Canada. Platform and Platform Computing refer to Platform Computing Corporation and each of its subsidiaries.090908