Platform™

IT Modernization: From Grid to Cloud in Financial Services

IT Modernization: From Grid to Cloud

Financial services (FS) institutions are still spending billions of dollars on IT, yet the current financial environment demands that these expenditures be curtailed in a radical way, by some estimates, budgets will fall up to 20%. Though the cut backs can't happen overnight, it is imperative that a solution is found to allow financial institutions to continue offering their services. Economic crisis or not, internal compute demand will continue to grow within most FS organizations. As the market consolidates and absorbs the disappearance or reorganization of some venerable organizations, there will be more business for fewer players. To overcome both of these obstacles, more demand for resources and less money, IT organizations will have to deliver more value through better utilization of existing infrastructure in order to drive ROI up. The technology that allows resource sharing is increasingly being viewed as a good solution. This technology is available today to share resources without making drastic cut backs to personnel and without impeding the organization's ability to continue innovating products and services.

Many FS companies are already sharing resources to maximize resource utilization in order to raise infrastructure ROI. IT departments are swiftly becoming service providers for LOBs, opening the door to opportunities for sharing common resources among a number of internal customers. The financial crisis is simply providing another reason to look at infrastructure options that enable FS companies to remain competitive and decrease the costs of IT infrastructure while making better use of the existing IT resources. In such an environment, cloud computing's economic resource sharing model enables on-demand IT resource delivery without additional capital expenditure.

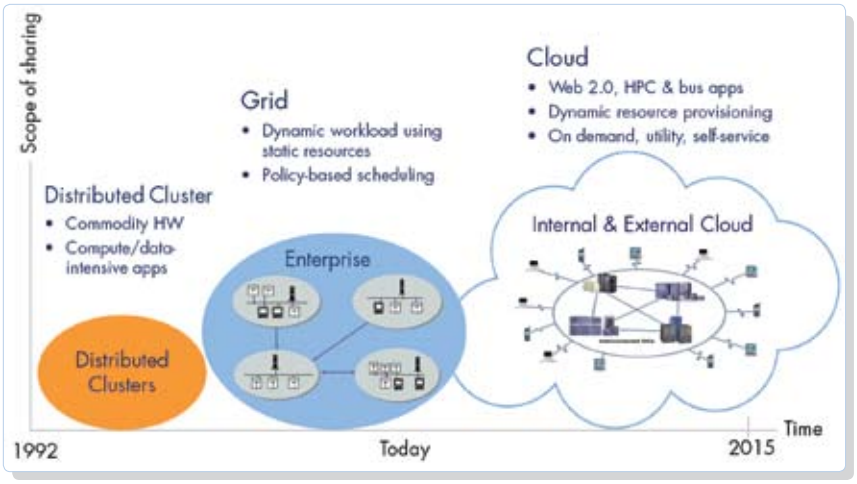

Figure 1: The evolution towards cloud computing

Cloud Computing: Reducing Data Center Waste

FS organizations are virtualizing IT resources to reduce costs and increase their ability to meet unpredictable current and future business requirements with a finite amount of resources. Today's data center is not structured to provide this sort of low-cost flexibility – the cloud model is.

Platform Computing is leading the way in leveraging cloud technology to help FS enterprises accelerate data center performance to the speed of business demand. By dynamically allocating resources to users on a real-time basis, removing bottlenecks and having better visibility into IT resources, organizations can increase resource utilization, defer the need for additional hardware acquisition, increase overall reliability of application performance, and meet organizational goals at the lowest total cost of ownership.

Cloud computing is the next evolution for data centers. A cloud is a data center architecture that automates the allocation and reallocation of IT resources to end-user applications in order to dynamically respond to changing business demands and priorities, while billing users for consumption. Companies deploying clouds benefit from cost savings, performance improvements, and enhanced operating flexibility. The promise of cloud computing is the ability to provision infrastructure quickly, on-demand, and to flexibly scale up or down based on fluctuations in actual demand and usage. Compute clouds can take many forms depending on the needs of the organizations deploying them. However they can be divided into two broad categories, "external clouds" offered by a service provider to paying customers or, "internal (private) clouds" deployed by the IT department of an enterprise to its internal users.

A hybrid cloud combines an internal cloud for mission critical, security-sensitive, or data-intensive application workload, with external cloud resources for less critical workloads in peak usage periods. Thus you can both minimize the cost of servicing higher SLA workloads by running in shared on-site infrastructure, and minimize the cost of servicing lower SLA workloads by dynamically leveraging the lowest-cost external provider. Many FS organizations are implementing a hybrid approach in their data centers. This approach provides a great degree of flexibility for IT service delivery, while minimizing costs through effective placement of applications according to SLA criteria.

Internal vs. External Cloud

Criteria to help determine what conditions must be met for a cloud model to succeed.

Internal Cloud

- Resources are underutilized
- Latency sensitive apps
- Apps with mission critical or highly customized SLAs
- Robust security requirements
- Dependencies on large data sets that change frequently
- Large IT operations

External Cloud

- IT requirements that are not well-defined
- Capital budgets are limited
- Applications with non-mission critical SLAs
- Security can be outsourced
- Data size is small and/or relatively static

Dynamic Resource Management: Optimizing Cloud Infrastructure

Dynamic resource management is the glue that maximizes the efficiency of cloud computing infrastructure, providing a solution for the barriers to internal clouds while making best use of external cloud capabilities. Dynamic allocation of resources is enabled through robust supply and demand policies that are automatically "matched" to achieve SLAs at lowest possible cost. Dynamic resource management is not about tearing the silo model apart. It is about finding new ways for IT to provide a sharing function that supports corporate goals, and engenders trust from application and budget owners. To deploy resources as needed, when needed, enterprises must have timely data about the state of their entire IT infrastructure. They need to be able to monitor their IT environment, manage the flow of application processes within the environment, detect, identify and route around anomalies, provide streamlined, secure access to resources for all users, anticipate changes in business demand for resources, and bill usage to consumers.

The most widely adopted dynamic resource management solution in the world, that does just this, is Platform Enterprise Grid Orchestrator (EGO). It is the enabling technology behind Platform Symphony and Platform LSF. Within a Platform-enabled internal cloud, enterprises can collaborate, share data and software, store increasingly large volumes of information, and reliably and securely access vast amounts of processing power – on demand. The internal cloud also allows effective alignment of the IT infrastructure with underlying business objectives for greater impact on corporate performance. But organizational barriers continue to pose obstacles to private clouds and shared computing – at enormous cost to the business. Organizations miss out on billions of dollars of potential savings each year when they fail to share IT resources across the business.

Enterprise-Class Dynamic Resource Management

Platform EGO (the kernel of all Platform products for years) is a robust, scalable, and trusted dynamic resource management product that enables IT managers to find new ways to provide a sharing function that supports corporate goals. Platform EGO's unique approach to sharing – leaving infrastructure silos and their corresponding ownerships alone – implements robust sharing policies outside the silos that grow and shrink available resources in a non-obtrusive manner. External cloud resources can be leveraged seamlessly when internal enterprise resources are insufficient, or where the cost model for external resources is appropriate.

Platform EGO's sharing policies control how cloud resources can be allocated among participating application owners. Priorities are set to ensure the highest priority applications have access to the right amount of resources to meet service levels. Once sharing policies are defined, Platform EGO dynamically places applications on appropriate resources to optimally utilize available supply. Since the impact to the business owner is a reduction in capital and operational expenditure, this hidden supply and demand matching engine quickly becomes a trusted component of the application's operations. Thus validated, the cloud business model can then expand to an economy of scale reflected in the bottom line.

Critical Elements of Cloud Infrastructure

Automation
Automation tools are a must in order to move logical resources between consumers, in a timely and guaranteed manner.

Virtualization
Virtualization enables pooling physical resources and deploying them in a flexible manner. Capacity can be added or removed on the fly, machines migrated to different physical locations, or rapidly repurposed.

Utility Infrastructure
Robust utility solutions, such as a common data fabric that spans data centers, enable enterprise applications to access data regardless where they're placed. These technologies are a significant requirement for transforming IT to run like a business.

Matching Supply with Demand
Meeting SLAs in a dynamic cloud infrastructure relies on dynamic resource management. Without it, cloud projects risk becoming nothing but science projects for non-production work.

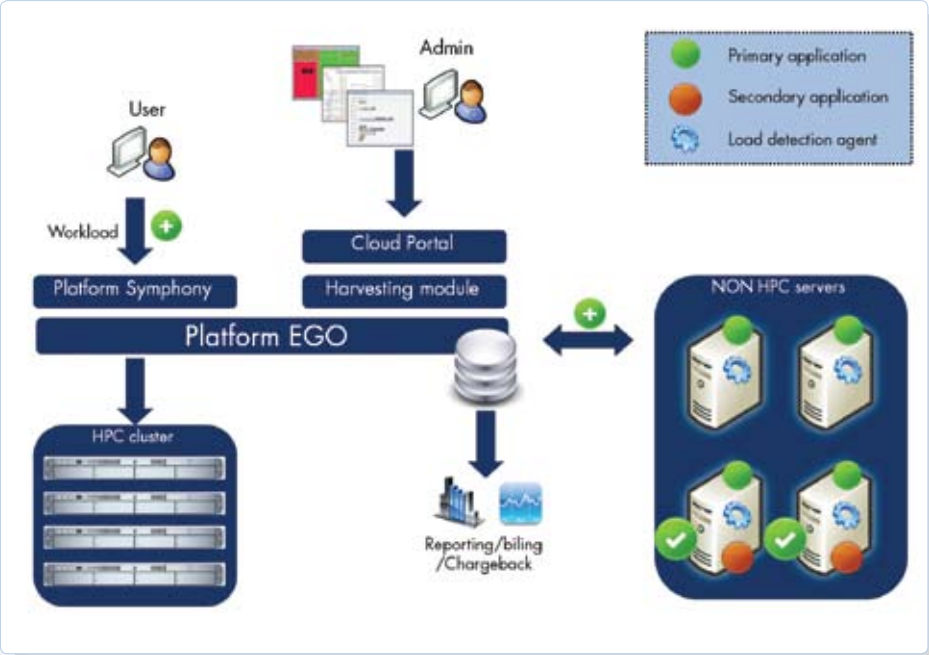

Figure 2: Build your internal cloud by using existing non-clustered machines to increase grid capacity on demand.

Platform EGO Sharing Models

The overwhelming resource sharing requirement from business owners is to retain control of the sharing process. Concurrently, business owners don't want to own the end-to-end operational management of a sharing solution; this would simply be exchanging capital expenses for operating costs. A centralized IT sharing broker is required to handle the day-to-day operations of shared infrastructure, while giving business owners the tools and authority they need to determine what resources and under what conditions sharing can occur. To enable this balancing act, Platform EGO provides two complementary sharing models that work in unison.

Peer-to-Peer Sharing – Individual application and resource owners define rules for when they will lend out resources and when they will borrow from others. Peer-to-peer trading leaves control of budgets within the business units yet facilitates improved utilization.

Centralized Sharing – A pool of resources is dynamically rented to application owners in the organization. Central IT managers define the rules for whom can get access to what, and when; the same model is employed by external cloud service providers. This sharing model is used to replace over-provisioning as an IT strategy. Instead of buying for peak demand, bursty usage patterns can be addressed by owning a smaller resource footprint, and supplementing it with additional resources only when they're needed.

Cloud deployments generally use a mixture of both sharing models. For this combined approach to be effective a centralized IT body is required to facilitate and broker the trading function. This IT brokerage provides a non-partisan, third-party operational function that breaks down the psychological barriers to sharing.

Optimizing virtualization

Platform EGO supports automated, policy-driven, virtual machine (VM) lifecycle management and self-service VM management needed for cloud projects. Several forms of virtualization are currently evolving that make it easier to pool and carve up physical hardware resources and deploy then dynamically. Servers and storage pools can be created through software abstractions and offered to users rather than having to physically procure, setup, and test hardware. Virtualization also enables mobility so that capacity can be added or removed on the fly, machines migrated to different physical locations, or rapidly repurposed. To maximize the ROI promised by virtualization, users need tools specifically built for managing enterprise virtualization deployments. Platform offers the only virtualization solution that meets the scale and dynamic computing requirements of enterprise data centers.

The bottom line

Enterprises cannot afford to waste their investments in IT by letting resources sit idle in one department while others are scrambling to add capacity. The answer is to enable sharing through the cloud. Dynamic resource management supports the cloud by enabling enterprises to build agile, responsive data center infrastructures that facilitate meeting strategic business goals in a fast-paced market environment. By matching current and future compute demand with available supply, as economically as possible, enterprises can turn their IT infrastructure from a static, inflexible cost, to an agile, strategic asset.

Platform EGO Benefits for Cloud Computing

Platform EGO provides the capabilities to truly maximize the amount of waste that can be shed by dynamically sharing IT resources.

Deliver IT Services Faster
Platform EGO automates resource allocation in response to changing supply and demand, enabling acceleration of resource delivery to users from weeks or months to minutes or hours.

Reduce IT Costs
Defer capital and operational expenses. Instead of procuring new machines to meet demand, resources are rented from application groups with low server utilization. The renter absorbs a smaller increase in operational expense – and no extra capital expenditure – while cross-charging nets the seller a decrease in operational expenses.

Increase Utilization, Guarantee SLA
Platform EGO eliminates manual demand and supply matching, for optimum use of cloud computing infrastructure. By sharing resources across different application groups and business units, FS organizations can increase utilization significantly. The policies that enable this ensure that application priorities are respected.

For a customer proof please read: "Virtualization Is Virtually Everywhere on Wall Street", Wall Street & Technology, Jan. 15, 2009, available on Platform's web site at http://www.platform.com/news/2009/virtualization-on-wall-street

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 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
Oslo: +44 1256 883756
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
lliew@platform.com

Copyright © 2009 Platform Computing Corporation. The symbols ® and T 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.051109