

Platform on track from grid to cloud; financial services will lead

Analyst: William Fellows

Sector: Enterprise Software

Sharing used to be a dirty word, but not any longer. Taking a utility model and applying it across an organization now seems prudent. Traditional investment bank high-performance computing (HPC) environments are being extended to support other applications and tasks such as retail and analytics across financial services industry (FSI) organizations, while self-service provisioning for test and development on virtualized infrastructure is a perfect application for internal cloud computing. The sharing of Web-server resources is behind this curve, mainly because of cultural reluctance to share resources. **Platform Computing's** work with existing customers suggests that savings of 30% are achievable by moving from a siloed deployment to an internal cloud.

The 451 Take

With a new tagline 'the power of sharing,' Platform is emerging from an HPC cocoon with a corporate initiative to double down on internal cloud-computing opportunities with early adopters, both in FSI and elsewhere. VMware integration is just the opening salvo in connecting different worlds. While it's not mainstream yet, Platform's clear belief is that future datacenters will look like internal clouds. We are seeing grid teams with successful services, especially those established within the investment-banking operations of financial services groups, being charged with taking the utility model and applying it more broadly across the group. The idea is that the benefits enjoyed by these groups should be able to support other user constituencies, services and workloads – subsequently offering a wider benefit across the organization. Federating and extending grids to support retail banking is a frequently mentioned example. Bank of America, Barclays, Citi, JP Morgan Chase, Merrill Lynch and Morgan Stanley are among those grid users we have looked at in our Grid Adoption Research Service.

Products

Platform's LSF, Symphony and VMO (Virtual Machine Orchestrator) products and the underlying EGO (Enterprise Grid Orchestrator) infrastructure are already supporting these kinds of activities, mostly for internal clouds, but also for external clouds, including **Alatum**, Singapore's national cloud. Other products using EGO as the base will expand the range of possible applications and workloads supported. EGO is the dynamic resource management function that has been unbundled from the products to match demand with supply across silos and, where required, to overflow pools or external resources.

Citi's use of LSF and Symphony has pushed utilization up from 20% to 80% across 20-plus intra and end-of-day pricing and risk applications. **JP Morgan Chase** has some 21

applications up on its compute backbone, which uses Symphony, and has reduced cost to \$0.56 per application hour.

The uncertainty that is still unfolding will drive enterprises to do more with existing infrastructure, which means sharing, with an emphasis on costs and ROI. Cloud computing is the logical endpoint for the convergence of grids, virtualization, automation and utility delivery models. Dynamically scheduled workloads can now have virtualized resources provisioned to meet their need – on demand. Grid nodes no longer need to be configured to a specific stack to handle only a specific set of jobs. The enabling technology here is Platform's VMO.

In a little-publicized move, VMO has recently been integrated with **VMware** (it also works with Xen and Kernel Virtual Machine), with Hyper-V integration to follow. Traditional job scheduling is dependent upon the operating-system image and application version residing on a given compute node. The integration of VMO with VMware, called Virtual Computing Cluster, dynamically changes the operating system running on a compute node (breaking the hard wiring of image to node) and manages how jobs are prioritized and controls which applications are instantiated on a node.

Many organizations – including large FSI companies – are talking to public cloud providers. However, the cost of service supplied by internal utilities continues to be competitive. The bottom line remains cost: whether an organization can do it more cheaply than a third-party provider. The \$0.50 per CPU/storage/network hour achieved by **Merrill Lynch** (see previous reports) is a benchmark that we've heard has been achieved at other end users too. We believe the sweet spot of activity in enterprises will continue to be on-premises or internal cloud computing. However, setting security and intellectual-property issues to one side for the time being (they will be figured out over time), the public cloud is well-positioned to offer a cheaper way of organizing IT because it moves capex spending to opex. FSI users will continue to look outside their organizations for resources on demand, and, ultimately, internal clouds will be augmented by outside services.

Beyond FSI

FSI users are early adopters of cloud computing, and the pace setters have derived enough benefit that they are driving others to it quickly. Indeed, moving from grid to cloud is increasing ROI for FSI companies.

The differentiator between FSI and other sectors is that in FSI there is no difference between HPC and non-HPC environments. They are all business applications. HPC is a part of enterprise architecture, and the sharing of applications at the enterprise level is a standard practice. In most other sectors, HPC exists in fiefdoms, usually unconnected with the rest of enterprise IT. In electronics and manufacturing, HPC has been mostly kept separate from virtualization and server-consolidation programs.

However, IT budgets look set to decrease by as much as 20%. Yet work still needs to get done, and computing demands will continue to grow. As IT is charged to increase ROI and make better use of existing resources, sharing will become more commonplace. IT departments will become service providers for lines of business.

Business model

Platform hasn't given away much financial information lately. It appeared to have stalled somewhere around the \$60m revenue mark on its goal to reach \$100m by 2010. But having doubled its revenue in the FSI sector to 15% of overall revenue in the year to March 2008 (see previous reports), we expect FSI to have grown incrementally again through this year. Platform claims 16 years of profitable and continuing growth overall, and will share additional details when its financial year ends in March. It claims 2,000 customers and has some 500 employees.

Platform made a tactical acquisition of **Scali** in 2008. Making HPC simpler and more manageable will continue to be a key element of its strategy in 2009. M&A will continue in this vein; Platform has already talked to a couple dozen cloud-computing startups.

A number of vendors have spoken about an increased opportunity for swapping out LSF enterprise customers to lower-priced alternatives as they reach the end of their enterprise-user license agreements. Platform says this has been addressed lately with an adjusted pricing model.

Competition

In FSI, Platform's key competitor is **DataSynapse**. EGO has remained a differentiator for it, although DataSynapse's upcoming product enhancements, including federation, will push Platform harder. In seeking the next set of opportunities after the first blush of enterprise grid deployments, the two companies went off in somewhat different directions. Platform, while retaining a high profile in FSI, had hunkered down in its core HPC markets while DataSynapse had a greater focus on service-oriented architecture.

Microsoft made a big splash for the launch of its Windows HPC Server 2008 last fall. It believes the operating system now has around two-thirds of the functionality of both Platform and DataSynapse products and is now a creditable swap-out opportunity at FSI organizations. It touts reference customers such as **Lloyds TSB Bank**.

VMware's Distributed Resource Scheduler makes it a rival to Platform's own scheduling tools (Platform was originally to have supplied scheduling to VMware); however, the two appear to have found new ground to cooperate with the arrival of cloud computing and the intersection of resource and workload management needed to support on-demand cloud computing inside organizations. Other HPC and grid middleware competitors include **Altair Engineering, Univa UD, Cluster Resources** and open source mechanisms such as Condor.

SWOT analysis

Strengths	Weaknesses
Whether it's called cloud computing or something else, resource sharing and policy-based scheduling will be a common theme for datacenters through 2009. Platform looks set to follow up its 2008 HPC acquisition with M&A in the cloud.	Cloud computing is riding into town on the back of an economic downturn with IT budgets being cut. Can grid companies break out of HPC fiefdoms to serve additional application and workload needs across multiple groups and different sectors in this climate? It's a big, but not impossible, reach.
Opportunities	Threats
The emergence of shared infrastructures or internal cloud as the logical endpoint for grid, virtualization, automation and utility strategies is an opportunity grid vendors have been struggling to find for some years now. It has taken the economic slowdown to move this into high gear.	It seems unlikely that in a year or two's time there can be as many different providers of distributed resource and workload scheduling/management products. The question is where they are rolled up as virtualization management, grid and incumbent systems management worlds collide.

About The 451 Group

The 451 Group is a technology industry analyst company focused on the business of enterprise IT innovation. The company's analysts provide critical and timely emerging-technology insight to clients at vendor, investor, services and end-user organizations – insight that aids both strategic and tactical decision making for competitive advantage.

The 451 Group is headquartered in New York, with offices in key locations, including San Francisco, London and Boston. The company also operates Tier 1 Research – an independent division of The 451 Group – which analyzes the financial and industry implications of developments impacting public and private companies within the IT, communications and Internet sectors.

For additional information on the company or to apply for trial access to its services, go to: www.the451group.com