Platform HPC Workgroup Manager

Simplified management of high performance computing clusters

A Platform Computing white paper
September 2009

Author:

Gord Sissons
Platform Computing
gsissons@platform.com

Platform Computing

Contents

1. Introduction

As any IT manager knows, cost and complexity often go hand in hand. There are few software environments more complex than modern high performance computing (HPC) clusters, so dealing with the problem of complexity in cluster management is central to containing cost.

Linux®-based commodity clusters have come to dominate supercomputing, and are making serious in-roads in corporate and academic settings as well. This is due in part to their superior performance, cost-effectiveness and flexibility.

Ironically, the same factors that make Linux® the clear choice for HPC professionals can make it less accessible to smaller centers that may have Microsoft Windows® administrators on staff but have little or no UNIX® or Linux experience. The complexity and associated cost of cluster management threatens to erode the very benefits that make open, commodity clusters so compelling. Not only can HPC cluster deployments be difficult, but the ongoing need to deal with heterogeneous hardware and operating systems, rapidly evolving toolsets, and changing HPC middleware combine to make deploying and managing an HPC cluster a daunting task. The industry needs a fresh approach.

This paper provides an overview of Platform HPC Workgroup Manager (HPC-WM). HPC-WM dramatically simplifies both the installation and ongoing management of Linux clusters, making them accessible to even the smallest organizations. It does this by offering a complete set of capabilities and ease of management features simply not found in other HPC cluster management solutions.

2. The cluster management challenge

To provide a proper HPC application environment, administrators need to deploy a full set of capabilities to their user communities as shown in Figure 1. These capabilities include Cluster Provisioning & Node Management, Application Workload Management, and an environment that makes it easy to develop, run and manage distributed parallel applications. Because modern application environments tend to be heterogeneous, some workloads will require Windows® compute hosts while others will require particular Linux operating systems or versions. The ability to change a node's operating system on-the-fly in response to changing application needs (referred to as adaptive scheduling) is also important since it allows HPC centers to maximize resource use and present what appears to be a larger resource pool to cluster users.

Cluster users increasingly demand web-based access to their HPC environment both for submitting workloads and for managing and monitoring their jobs once submitted. An easy to use application-centric web interface can have tangible benefits including improved productivity, reduced training requirements, reduced errors rates, and remote secure access.

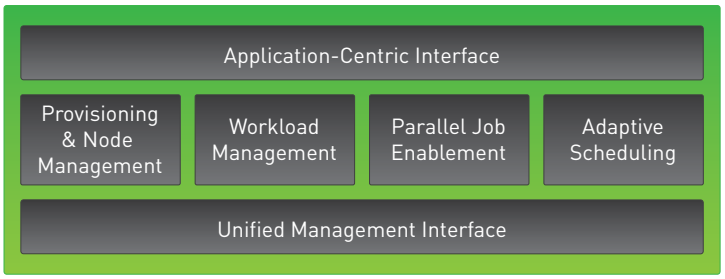

Figure 1. Essential components of an HPC cluster solution

[1]MPI refers to the message-passing interface, a widely deployed solution in HPC environments that enables distributed parallel applications by facilitating communications between processes running on different machines.

[2]OFED refers to the OpenFabrics Enterprise Distribution

[3]Anaconda is the installation program used by Fedora, Red Hat Enterprise Linux and other Linux Distributions

While there are several cluster management tools that address parts of these requirements, few address them fully, and some tools are little more than collections of discrete open-source software components.

Cluster toolkits such as NPACI ROCKS, IBM® xCAT and Scyld Clusterware™ focus largely on the problem of cluster node provisioning and management. While they clearly simplify cluster deployment, administrators wanting to make changes to node configurations or customize their environment will quickly find themselves hand-editing XML configuration files or writing their own shell scripts. Third-party workload managers and various open-source MPI[1] implementations might be included as part of a distribution, but these included components are loosely integrated at best, and they often represent a separate point of management and an entirely new subsystem that needs to be learned by a cluster administrator to be managed effectively.

Other HPC solutions are oriented to the particular problem of application workload management. These include solutions such as Condor®, Sun Grid Engine (SGE), MOAB Cluster Suite®, and PBS Professional™. While these are all capable workload managers, most do not address at all the issue of cluster node management, application-centric portals or adaptive scheduling. If such capabilities exist at all they usually require the purchase of additional software products.

Parallel job management is critical as well. The whole reason that customers deploy HPC clusters is to maximize application performance, and processing problems in parallel is a common way to achieve performance gains. The choice of MPI, its scalability, and the degree to which it is integrated with various OFED[2] drivers and high performance interconnects has a direct impact on delivered application performance. Furthermore, the workload manager needs to incorporate specific parallel job management features, or busy cluster users and administrators can find

themselves spending time cleaning up after failed MPI jobs or writing their own shell scripts to do the same.

As with any IT environment, on-going management represents the largest single cost associated with HPC clusters. If a cluster management solution is not complete and fully integrated, customers will quickly find themselves in the technology integration business, needing to deal with the problem of "what works with what" and grappling with a plethora of discrete software components, each having different management interfaces and software pre-requisites. Upgrading software in such an environment becomes a major challenge, since there is no guarantee that the newly added software component will work with other components in the HPC software stack.

Complexity is a real problem. You can imagine a small organization or department grappling with a new vocabulary of cryptic commands, configuring and troubleshooting Anaconda[3] kick start scripts, finding the correct OFED drivers for specialized hardware, and configuring open source monitoring systems like Ganglia or Nagios®. Without an integrated solution, administrators need to deal with dozens of distinct software components, and progress with HPC cluster implementations can slow to a crawl.

3. Re-thinking HPC clusters

Clearly these challenges demand a fresh approach to HPC cluster management. Platform HPC Workgroup Manager (HPC-WM) represents a new approach to how open HPC clusters are deployed and managed.

Rather than addressing only part of the HPC management puzzle, HPC-WM addresses all facets of cluster management including:

- Cluster provisioning and node management
- Workload management
- Parallel application enablement
- Adaptive scheduling
- Unified management
- Application-centric user interfaces

Optimized specifically for clusters of up to 32 nodes, HPC-WM includes Platform Cluster Manager (PCM), an easy to use web-based cluster provisioning and management tool that offers unique capabilities not found in other cluster management tools. Industry leaders including Dell™, Intel®, HP and Red Hat® have all embraced PCM as the basis of their own HPC cluster management solutions, recognizing its flexibility and ease of management features.

Platform LSF Workgroup Edition, Platform Computing's popular workload management solution is also included as a standard part of HPC-WM. Unlike other solutions where the scheduler is an add-on that needs to be separately installed and configured, Platform LSF is pre-configured and ready for use on installation. Platform LSF is integrated directly into the Platform Management Console (the web-based interface to HPC-WM) and users can immediately begin submitting and monitoring their jobs through a configurable web-based job submission and management interface.

In addition to various industry standard MPIs, HPC-WM also includes Platform MPI, a high-performance MPI implementation that delivers superior application performance owing to its core-aware multi-threaded architecture along with ease of management features and parallel application debugging capabilities.

Also included with HPC-WM, and integrated into the web-based Platform Management Console interface, is adaptive scheduling achieved through an integrated "multi-boot" capability. The adaptive scheduling facility is transparent to cluster users and detects the resource requirements of user-submitted workloads and changes the topology of the cluster accordingly providing an adaptive infrastructure. This facility reboots nodes as necessary, subject to administrator defined policies, to achieve an optimal mix of operating system types on cluster nodes depending on changing workload requirements.

Other components and capabilities included in Platform HPC Workgroup Manager are:

- Collections of the latest cluster management tools from the open-source community
- An automated software maintenance and patching facility (via RHN or Yum[4])
- Pre-integrated HPC tools, libraries and developer tools
- An Intel® Cluster Ready[5] certified software suite
- An easy to use web interface used for cluster management
- A web portal that makes it easy to deploy self documenting application specific interfaces
- A full NVIDIA® CUDA environment to streamline the deployment of GPU-based clusters
- 24x7 enterprise support, backed by Platform and leading software providers and OEMs

The components that comprise Platform HPC Workgroup Manager are shown in Figure 2.

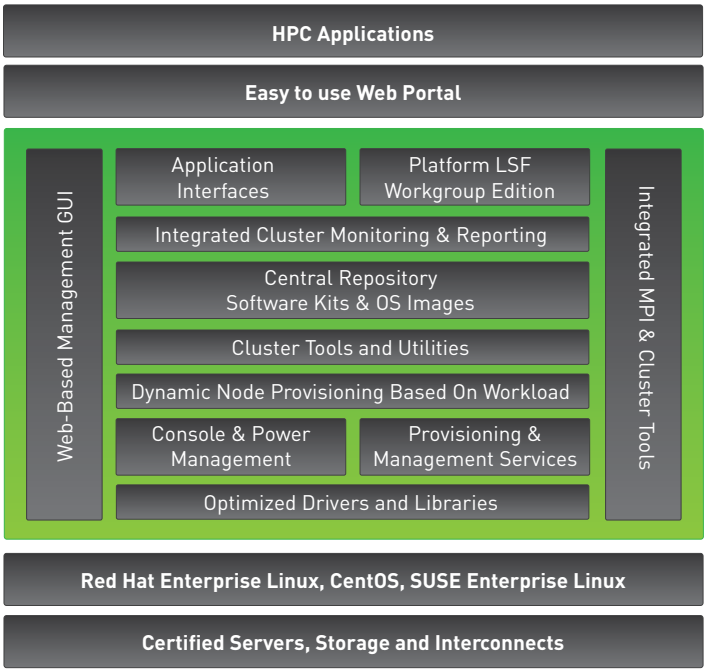

Figure 2. Platform HPC Workgroup Manager components

[4]Yum (yellow dog update manager) originally developed by Yellowdog Linux is a tool included in major Linux distros for automating package maintenance for systems that use RPMs for distributing packaged tools and applications

[5]The Intel® Cluster Ready (ICR) program and technology package makes it easy to design, build and deploy clusters built with Intel components.

4. Unique features

Platform HPC Workgroup Manager provides a wealth of features not found in other HPC cluster management solutions. Earlier we described the key capabilities required in an HPC cluster management solution. In this section we'll re-visit each of these functional areas and describe the unique features and benefits of HPC-WM.

4.1 Unified management interface

HPC-WM features a unified management console that is used to administer all aspects of the cluster environment. Through the management console, administrators can install and manage cluster nodes and interact with interfaces for cluster monitoring, reporting, charting and workload management.

As changes are made to the cluster configuration through the management console, HPC-WM automatically re-configures key components as illustrated in Figure 3. Because configuration of the cluster is automated, administrators don't actually need to know how to install, configure or administer the various included tools like Ganglia, ntop, and Nagios®. This is because these components are configured and re-started as required automatically whenever the configuration of the cluster is changed. Tools included in HPC-WM become immediately ready for use through the web-based interface. This degree of automation saves the cluster administrator time, means they don't need to become expert in the administration of open-source software components, and it reduces the possibility of errors and time lost due to mis-configuration. Cluster administrators enjoy the best of both worlds – easy access to powerful web-based cluster management tools without the need to learn and separately administer all the tools that comprise the HPC cluster environment.

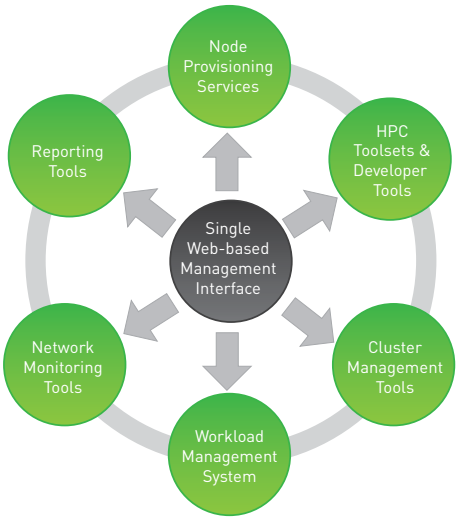

Figure 3. Changes in the web interface result in automated re-configuration of components

The Platform Management Console is pictured below. Whereas competing cluster management tools either have no web-based interface or multiple different interfaces for managing different functional areas HPC-WM presents a single unified interface through which all administrative tasks can be performed including node-management, job-management, jobs and cluster monitoring and reporting. Using the management console, even cluster administrators with very little Linux experience, can competently manage a state of the art HPC cluster.

Figure 4. The PMC provides a single interface for cluster management

4.2 Cluster provisioning and management

Fast and efficient software Installation – Platform HPC Workgroup Manager can be installed and configured in under one hour on a master node by inserting an installation DVD and answering a series of simple questions. Layered software packages are deployed as self-describing "kits", and installing software on cluster nodes is simply a matter of associating cluster nodes with flexible node group definitions through the web-based interface.

Flexible provisioning – Platform HPC Workgroup Manager offers multiple options for provisioning Linux operating environments that include:

- Package-based provisioning (SUSE or RHEL)
- Image based provisioning
- Diskless node provisioning

Large collections of hosts can be easily added as a group, and HPC-WM attends to details like IP address assignment and node naming conventions that can reflect the position of cluster nodes in data center racks. Unlike competing solutions, HPC-WM can deploy multiple operating systems and OS versions to a cluster simultaneously including Red Hat Enterprise Linux, CentOS and SUSE Linux Enterprise Server. This provides administrators with greater flexibility in how they serve their user communities, and means that HPC clusters can grow and evolve incrementally as requirements change.

Repositories, kits and components – Kits form the basic unit of software distribution in HPC-WM. Kits understand their own software dependencies and pre-requisites, and can install, configure, and un-install themselves. Many kits employ "plug-ins" that allow included software components to re-configure themselves automatically in response to external events such as a change in the configuration of the cluster or the personality of a node.

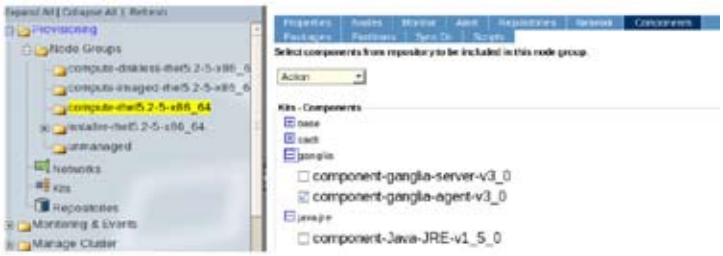

Figure 5. Software components are install by including the component in a node group definition

Tedious tasks such as upgrades and patches can be automated using standard facilities such as Yum and HPC-WM integrates with the Red Hat Network (RHN) service for those with RHN subscriptions.

Pre-configured software kits – Among the standard software kits provided for Platform HPC Workgroup Manager are:

- **Base kit** – Contains all the tools and applications required to manage the cluster
- **Cacti kit** – An open source reporting tool used to collect and graph various node metrics
- **Ganglia kit** – An open-source resource monitoring solution
- **HPC kit** – A collection of tools, MPIs, libraries and utilities
- **Platform LSF kit** – Platform LSF Workgroup Edition
- **Platform MPI kit** – An optimized MPI providing superior performance and management
- **Nagios kit** – An open source host, services and network monitoring solution (NMS)
- **NTOP kit** – A tool to monitor network bandwidth and analyze traffic
- **OFED kit** – A collection of drivers supporting server and storage interconnects and facilities such as IP over InfiniBand (IPoIB)
- **PCM GUI** – The Platform Management Console

- **CUDA kit** – The full CUDA® development and run-time environment making it easy to deploy and manage HPC clusters exploiting the power of NVIDIA® Tesla® GPUs.

HPC-WM can support a practically unlimited number of node group definitions that incorporate different operating systems and different layered software tools. Changing the configuration of a node is simply a matter of assigning that node to a different node group. Installing or de-installing software components is accomplished by clicking a check box beside the components to be installed on the node group's member hosts as shown in Figure 5. HPC-WM supports multiple repositories, based on different operating systems and versions, and node groups draw software components from these different repositories providing unmatched flexibility.

Cluster File Manager – HPC-WM includes CFM (Cluster File Manager) and PDSH (a parallel distributed shell). The Cluster File Manager simplifies administration and increases cluster availability by allowing changes such as new package installations, patch updates, and changes to configuration files to be propagated to cluster nodes automatically without the need to re-install cluster nodes. PDSH as its name implies is a mechanism that allows experienced administrators to perform operations quickly in parallel across multiple cluster nodes.

Integrated cluster monitoring and reporting – With HPC-WM, reporting and monitoring functions are integrated directly into the Platform Management Console. Monitoring facilities include node monitoring, service monitoring, graphing capabilities, workload monitoring and network monitoring. Various standard reports are produced automatically including cluster and host availability reports, workload and utilization reports and inventory summaries.

Repository snapshots / trial installations – Upgrading software can be risky, particularly in complex environments. If a new software upgrade introduces problems, administrators often need to rapidly "roll back" to a known good state. With other cluster managers this can mean having to re-install the entire cluster. HPC-WM incorporates the notion of repository snapshots, in essence a "restore point" for the entire cluster. Administrators can snapshot a known good repository, make changes to their environment, and easily revert to a previous "known good" repository at any time in the event of an unforeseen problem. This powerful capability takes much of the risk out of cluster upgrades.

Support for new hardware models – Over time, cluster administrators will likely want to add new hardware to their clusters, and this new hardware may require new or updated device drivers not supported by the OS environment on the installer node. This means that a newly updated node may not be able to network boot and provision itself until the installer node on the cluster is updated with a new operating system, a tedious and disruptive process. HPC-WM includes a driver patching utility that allows updated device drivers to be "poked" into existing repositories, essentially "future-proofing" the cluster and provide a simplified means of supporting new hardware without needing to re-install the environment from scratch.

Software updates with no re-boot – Some cluster managers take the approach of always re-installing nodes, regardless of how minor the change in configuration is. This is simple of course, but on busy clusters scheduling downtime can be difficult and disruptive. HPC-WM performs updates intelligently and selectively via the Cluster File Manager so that compute nodes continue to run even as non-intrusive updates are applied. The repository is updated as well so that future re-installs will reflect the same changes made "live" on cluster nodes. For changes that do

require the re-installation of the node (changing an operating system for example) these changes can be made "pending" until downtime can be scheduled.

4.3 Integrated workload management

Integrated workload management – While other cluster management solutions require that a workload manager to be separately sourced, installed and configured, HPC-WM includes Platform LSF Workgroup Edition. Optimized for simplified deployment and management, Platform LSF Workgroup Edition retains all of the powerful features of Platform LSF, but is fully configured on installation automatically, and is immediately accessible to both cluster administrators and users through the management interface.

Figure 6. Users can easily manage their own Platform LSF jobs

Application centric portal – Application specific job submission templates for popular HPC tools are included in the HPC-WM. By providing self-documenting, user friendly interfaces, training requirements and submission errors are reduced. These application centric interfaces are ideal for organizations supporting remote user communities, or educational institutions where HPC tools may need to be deployed to student populations in a controlled fashion.

4.4 Parallel job management

Platform MPI – The HPC kit in HPC-WM includes various industry standard MPI implementations including MPICH1, MPICH2 and MVAPICH1 optimized for cluster hosts connected via InfiniBand, iWARP or other RDMA based interconnects. Unique to HPC-WM is Platform MPI (formerly Scali MPI), a value added MPI implementation that provides superior management features along with better performance and scalability. Optimized for Platform LSF, Platform MPI is specifically optimized for multi-core systems. In the industry standard SPEC MPI 2007 benchmark suite, Platform MPI exhibits up to 42% better performance than its competitor on the Fluid Dynamics test component, and 17% better performance overall. A 17% improvement is dramatic in HPC environments where seconds count. Platform MPI an essential tool for any HPC environment running distributed parallel workloads.

4.5 Adaptive scheduling

Adaptive scheduling & dynamic provisioning – With its multi-boot manager, cluster hosts can re-boot themselves from a different disk partition automatically in response to changing application workloads. Through the management console, administrators can choose to deploy a reasonable mix of Windows® and various Linux environments depending on the anticipated needs of their user communities. As users submit jobs or run applications that have particular operating system dependencies (expressed using the graphical portal interface or via the Platform LSF command line), the composition of the cluster can "flex" within configurable limits as shown in Figure 7. Rather than being an extra-cost item as it is with other HPC management suites, this capability is included as a core feature of the HPC-WM.

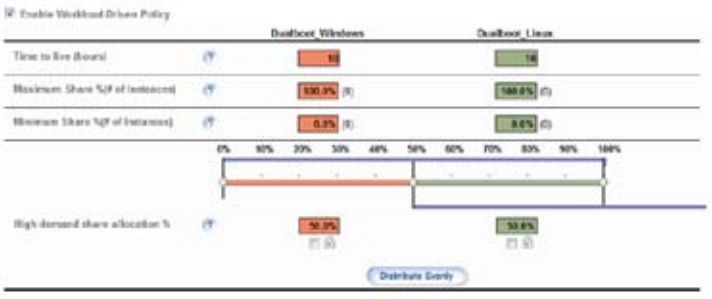

Figure 7. The cluster composition can "flex" as workload demands change

4.6 Service and support

Certified cluster configurations – HPC-WM is specifically tested and certified on partner hardware platforms. By qualifying each platform individually, and by providing vendor specific kits with optimized libraries and drivers that take maximum advantage of unique hardware features, Platform Computing has essentially done the integration work in advance, allowing clusters to be deployed quickly and predictably with minimal effort. HPC-WM is also certified with the Intel Cluster Ready program.

Enterprise class service and support – Widely regarded as having the best HPC support organization in the business, Platform is uniquely able to support an integrated HPC platform. This is because Platform Computing develops all of the key software components included in HPC-WM including the cluster provisioning and management tool, the workload manager, the GUI, MPI, the portal and the adaptive multi-boot facilities. Because support personnel have direct access to the developers of all software components, Platform Computing is able to offer a higher level of support and ensure that any problems encountered are resolved quickly and efficiently.

5. A complete solution

Compared to other cluster management and workload management solutions, Platform HPC Workgroup Manager simply provides the richest set of capabilities without compromising on the openness that makes Linux cluster so attractive to HPC users. This is because HPC-WM provides an integrated solution that addresses all aspects of cluster management including provisioning, node management, workload management, reporting and parallel job management. Table 1 below contrasts some of the major features of HPC-WM with those other cluster and workload management solutions which address only some parts of HPC management problem.

Comparison with other offerings			
Capability / Feature	**Platform HPC Workgroup Manager**	**An Open Source Tool**	**A Commerical Tool**
Cluster Provisioning and Management Features			
Initial cluster provisioning	✔	✔	*
Multiple provisioning methods	✔	✘	✘
Web-based cluster mgmt.	✔	✘	✔
Node updates with no re-boot	✔	✘	✘
Repository snapshots	✔	✘	✘
Flexible node templates	✔	Weak	*
Multiple OS and OS versions	✔	Weak	*
Workload Management and Application Integration Features			
Integrated workload management	✔	✔	✔
Application-centric web-portal	✔	✘	✘
HPC libraries and toolsets	✔	✔	✘
NVIDIA CUDA SDK support	✔	✔	*
Web-based job management	✔	✘	✔
Multi-boot based on workload	✔	✘	*
Advanced parallel job management	✔	✘	✔
Commercial grade MPI	✔	✘	✘

Table 1. Capability by product or suite

** Items marked with an asterix indicate that the feature is available but requires a separately installed add-on package.*

6. Sources of cost savings

Employing Platform HPC Workgroup Manager can help save or avoid cost in a number of areas. Some of the specific areas of cost savings are identified below.

Personnel savings – HPC-WM can substantially reduce the amount of effort required to deploy and manage a cluster. For even a moderately experienced system administrator building a fully functional HPC cluster and assembling and configuring all of the various layered software components can easily take weeks. HPC-WM can reduce this effort to a matter of a few days and can be effectively installed and managed by non-specialists.

Time is money – HPC-WM ensures that clusters are immediately productive, and that cluster downtime is kept to an absolute minimum owing to features like the CFM and repository snapshots.

Hidden cost avoidance – Some organizations account for additional system and network management tools outside of the context of their HPC deployment. However, these costs are real, and the features included standard in HPC-WM can come at an additional cost with other cluster management solutions. Examples of extra components include things like workload management systems, the cost of building or implementing a portal, dynamic provisioning and management tools and the cost of technical support.

Reduced error rates – A commonly overlooked source of cost and reduced productivity is related to user errors in job submission. A 16-way MPI job that runs for two hours, that needs to be run again because the user submitted the job incorrectly, costs 32 hours of CPU time. During this period, other jobs may need to pend in the workload management system waiting for available resources. With features like application-centric portals, and by allowing users to monitor their jobs in real-time, the opportunity for job submission errors are reduced, and if there is an error, users can recognize the error quickly and take corrective action without bothering an administrator.

Improved cluster utilization – There is a significant difference in cost and productivity between a cluster running at 80% utilization versus a cluster running at 95% utilization. Platform LSF Workgroup Edition, included in Platform HPC Workgroup Manager is widely regarded as among the best commercial workload management systems. By taking advantage of the sophisticated scheduling capabilities of Platform LSF, customers can not only better align cluster resources to business needs, but they can achieve better levels of utilization as well, effectively doing more with less.

Superior price / performance – Because the supplied tools and libraries in HPC-WM are optimized for selected vendor hardware configurations, customers are assured that they are getting the highest level of performance possible from their clusters. With a "do it yourself" approach to building clusters, customers risk deploying unoptimized libraries and tools which may work, but at reduced performance levels. As cited earlier, based on the SPEC MPI benchmarks, application benchmarks including fluid mechanics, quantum chromodynamics, heat transfer, weather forecasting and hydrodynamics all ran a minimum of 20% faster using the optimized Platform MPI included in HPC-WM than on other MPIs running on identical hardware.

Reduce risk – By relying on pre-integrated, pre-tested, certified configurations fully backed by Platform Computing, customers are assured that any issues can be addressed quickly and efficiently without the need for on-site consultants or additional support expertise that might result in unbudgeted costs.

Infrastructure cost avoidance – Many sites run mixed Linux and Windows environments. Rather than provision hardware to peak requirements for applications in each environment, total costs can be reduced by automatically changing the personality of nodes dynamically in response to application demand. With its built-in multi-boot manager HPC-WM includes this adaptive scheduling capability as a core part of the solution.

"Future-proofing" – As explained earlier, management costs tend to dominate overall IT related spending. A cluster might run for a few years, but when it comes time to add nodes to that cluster, upgrades can be a challenge since new platforms generally require new operating systems. HPC-WM provides a number of features that specifically "future proof" clusters including the ability to easily add device drivers accommodating new hardware without the need to re-install the head-node. Also, it allows different hardware types to draw their configuration from different repositories based on different OS versions. This effectively means that clusters can easily be grown incrementally avoiding costly downstream "rip and replace" scenarios. As administrators become confident that the new hardware is functioning properly, they can easily migrate newer cluster nodes to production node groups and similarly retire older nodes on the same cluster.

7. Summary

Platform HPC Workgroup Manager is the ideal solution for deploying and managing state of the art HPC clusters. It makes cluster management simple, enabling analysts, engineers and scientists from organizations of any size to exploit the power of open Linux clusters with a pre-integrated, vendor certified software solution.

Unlike other HPC solutions that address only parts of the HPC management challenge, HPC-WM uniquely addresses all aspects of cluster and management including:

- Cluster provisioning and management
- Workload management
- Parallel job enablement
- Adaptive scheduling
- Intuitive application-centric interfaces
- A unified management interface
- A single source of quality support

By specifically focusing on simplifying management over the entire life-cycle of the cluster, Platform HPC Workgroup Manager has a direct and positive impact on productivity while helping to reduce complexity and cost. The comprehensive web-based management interface, and features like repository snapshots and the ability to update software packages on the fly means that state of the art HPC clusters can be provisioned and managed even by administrators with little or no Linux administration experience.

With Platform HPC Workgroup Manager, administrators can spend less time worrying about cluster management, and more time focusing on their applications and the needs of their user communities. Similarly, cluster users can spend more time focused on their own projects and research.

Platform Computing

Platform Computing is the leader in cluster, grid and cloud management software - serving more than 2,000 of the world's most demanding organizations for over 17 years. Our workload and resource management solutions deliver IT responsiveness and lower costs for enterprise and HPC applications. Platform has strategic relationships with Cray, Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat®, and SAS®. Visit www.platform.com.

World Headquarters
Platform Computing Corporation
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

Partner Sales
Toll-free Tel: 1 877 710 4522
partnersales@platform.com

North America
New York: +1 646 290 5070
San Jose: +1 408 392 4900

Europe
Bramley: +44 (0) 1256 883756
London: +44 (0) 20 3206 1470
Paris: +33 (0) 1 41 10 09 20
Düsseldorf: +49 2102 61039 0
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 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.120809