

Using & Managing HPC Systems



Presenter: William Lu, Ph.D., Marketing Director
Date: October 7, 2008

Using and managing HPC systems, organizations are facing some challenges:

Complex hardware

- Performance relies on CPU, memory, storage, interconnects,...

Reliability

- Commodity hardware has good performance/cost, but...

Users competing resources

- Why my jobs are waiting?

METHODOLOGY

The sources for this white paper on the origins, recent developments, and importance of HPC management software are as follows:

all term
sponsi
hardwa
conduc
Devices
(SAS) to

IN THE

This white
software
and cons
elaborate

This white
out the el
operating
suggests w
not. It poin
on a path
hardware a
with manag
leading prov
customers a

SITUATION

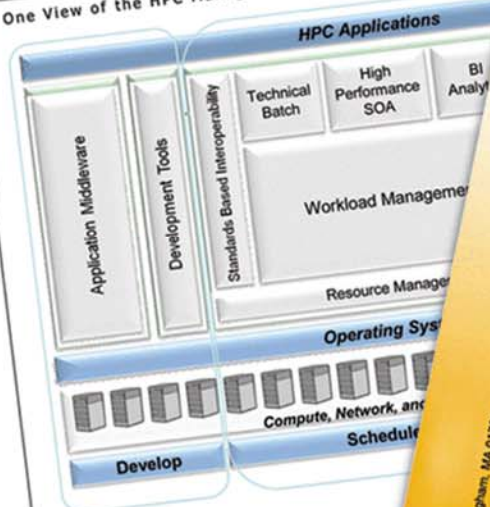
From its origi
processing (T
transactions p
applications a
diverse role. F
of crucial linki

HPC Management
Middleware

By virtue of this
software compo
system as midd
the many distin
working between
shortcoming by
software presents

FIGURE 1

One View of the HPC Management Software Stack



Source: Platform Computing, 2008

Is HPC management software middleware? Yes, because the application and the OS are the two main components of the HPC management software stack. The growing complexity of parallel HPC applications and the growing complexity of parallel HPC management software have led to the development of HPC management software. The anatomy of the HPC management software stack is carefully articulated to be well understood. The proper identification of new components is a key to the evolving infrastructure underlying HPC. The further transform HPC and become the key to today's HPC management software.



WHITE PAPER

HPC Management Software:
Reducing the Complexity of HPC Cluster and Grid
Resources

Sponsored by Platform Computing

Steve Conway

Earl C. Joseph, Ph.D.

May 2008

Richard Walsh

IDC OPINION

The typical architecture of a high-performance computing (HPC) resource is no longer a single custom-engineered, integrated hardware and software system with collections of vendor-engineered, integrated hardware and software components typically constructed à la carte from standards-based hardware and software component technologies. These resources are likely to include multiple systems at multiple sites. The variety and number of those components, their distinct providers, their potentially different locations, and the costliness of testing user applications against them all have driven up HPC resource management complexity and created an integration gap. Most HPC system configurations purchased today are unique assemblages of hardware and software. The growing task of integrating them into a reliable, productive, and secure working environment has fallen heavily on system administration and user support personnel.

The careful design, selection, and integration of the software components between the operating system (OS) and application layers of the software stack are increasingly vital for reducing the complexity of HPC cluster and grid management and for presenting an integrated and responsive HPC resource to users. Often informally referred to as *middleware*, the number of components in midstack HPC management software has grown significantly in the past 10 years. In IDC's opinion, this process has exposed a growing opportunity for vendors to add value to the HPC work cycle by adopting, streamlining, and modularizing the now long list of software functions assumed by HPC management software, including:

- ☑ User application compilation, debugging, and profiling
- ☑ User and system application input/output translation and scripting
- ☑ User application libraries, mathematical, IO, and parallel (e.g., MPI, OpenMP)
- ☑ Workload queuing, scheduling, and management (e.g., migration, checkpoint, restart)
- ☑ System and application installation, integration, and patching
- ☑ System management — servers, software, security, disk storage, and backups
- ☑ System and application monitoring, reporting, provisioning, reconfiguration, and failure detection

The software
between the
application and
the operating
system...

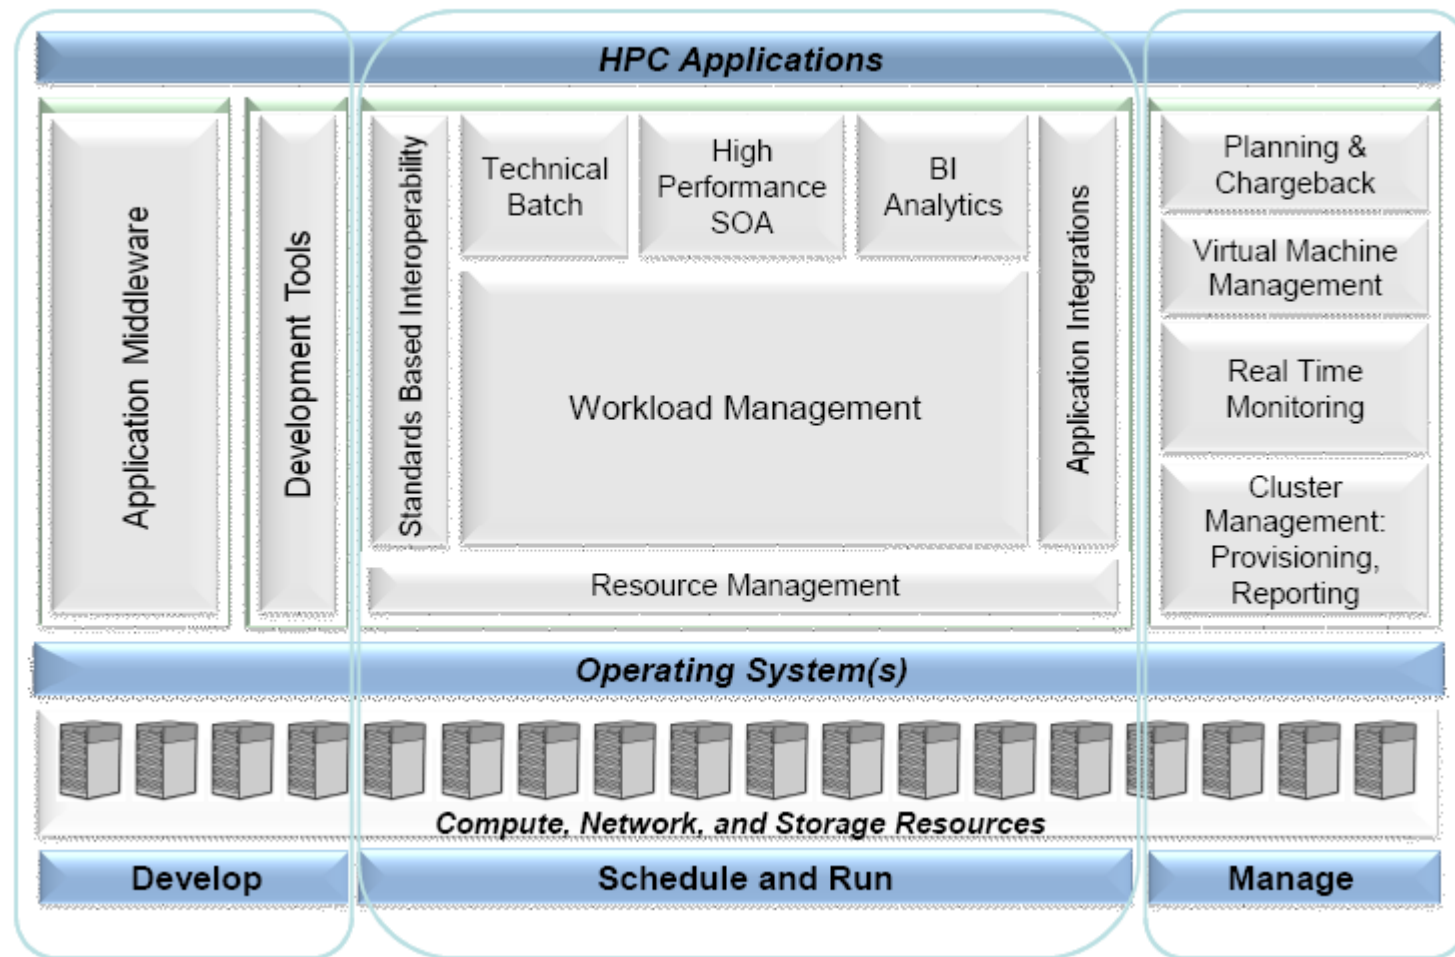
Applications

HPC Management Software

Hardware and OS

... that enables users to develop, run
and manage compute or data intensive
applications and manage HPC systems.

One View of the HPC Management Software Stack



4,000,000

Managed CPUs

2,000

Customers worldwide

500

Employees in 15 offices

16

Years of profitable growth

1

Leader in HPC

Gov, Research & Edu

- CERN
- DoD, US
- DoE, US
- ENEA
- Georgia Tech
- Harvard Medical School
- Japan Atomic Energy Inst.
- MaxPlanck Inst.
- MIT
- SSC, China
- Stanford Medical
- TACC
- U. Of Georgia
- U. Tokyo
- Washington U.

Financial Services

- BNP
- Citigroup
- Fortis
- HSBC
- KBC Financial
- JPMC
- Lehman Brothers
- LBBW
- Mass Mutual
- MUFG
- Nomura
- Prudential
- Sal. Oppenheim
- Société Générale

Industrial Mfg.

- Airbus
- BAE Systems
- Boeing
- Bombardier
- Deere & Company
- Ericsson
- Honda
- General Electric
- General Motors
- Goodrich
- Lockheed Martin
- Nissan
- Northrop Grumman
- Pratt & Whitney
- Toyota
- Volkswagen

Oil & Gas

- Agip
- BP
- British Gas
- China Petroleum
- ConocoPhillips
- EMGS
- Gaz de France
- Hess
- Kuwait Oil
- PetroBras
- Petro Canada
- PetroChina
- Shell
- StatoilHydro
- Total
- Woodside

Electronics

- AMD
- ARM
- Broadcom
- Cadence
- Cisco
- Infineon
- MediaTek
- Motorola
- NVidia
- Qualcomm
- Samsung
- Sony
- ST Micro
- Synopsys
- TI
- Toshiba

Life Sciences

- Abott Labs
- AstraZeneca
- Celera
- DuPont
- Eli Lilly
- Johnson & Johnson
- Merck
- National Institutes of Health
- Novartis
- Partners Health Network
- Pharsight
- Pfizer
- Sanger Institute

Other Industries

AT&T

Bell Canada

Cingular

DreamWorks Animation SKG

GE

IRI

Telecom Italia

Telefonica

Walt Disney Co.



Platform OCS 5 and Platform Manager integrated in Dell cluster systems



Platform LSF, Platform Manager form key parts of Unified Cluster Portfolio



Platform enterprise solutions support a wide range of IBM HPC systems



Platform delivers first certified Intel® Cluster Ready solution, Platform OCS 5



Integrates Platform LSF and Platform Symphony in grid solutions



Platform OCS 5 powers the Red Hat® HPC Solution



OEMs Platform's core technology in SAS® applications



Develop

Schedule & Run

Manage

Platform Accelerate

Platform MPI

Platform LSF

Platform LSF MultiCluster

Platform LSF License Scheduler

Platform LSF Session Scheduler

Platform Process Manager

Platform EGO

EnginFrame

Platform Manage

Platform Manager

Platform Analytics


Platform RTM

Platform VM Orchestrator

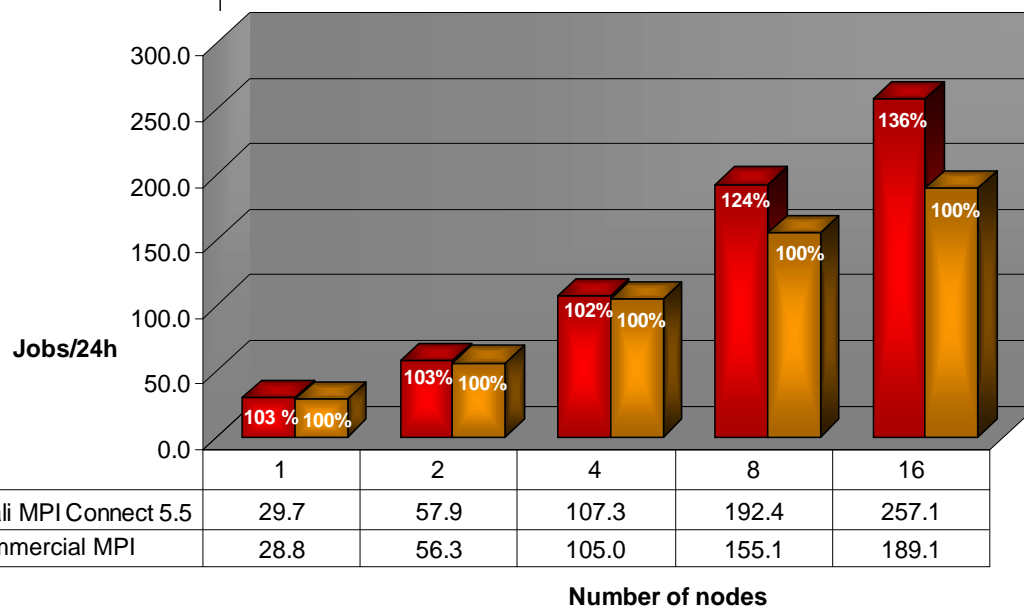

Platform Open Cluster Stack 5: Accelerate & manage single Linux cluster



- Platform MPI contains *forward looking* MPI features:
 - Parallel Checkpoint/Restart
 - Dynamic affinity (load monitoring)
 - Windows CCS support (OS agnostic clusters)
 - Dynamic process creation
 - Network Failover
 - Replace run time (.so file) without recompile/relink




Platform MPI out-performs by ~30%




- Robust, fault tolerance, and high performance
- Powering mission critical environment
- Special price for education

By scheduling workloads intelligently according to policy, Platform LSF reduces application run-time and optimizes resource use.



VIRTUALIZED VIEW OF COMPUTE, NETWORK AND STORAGE RESOURCES



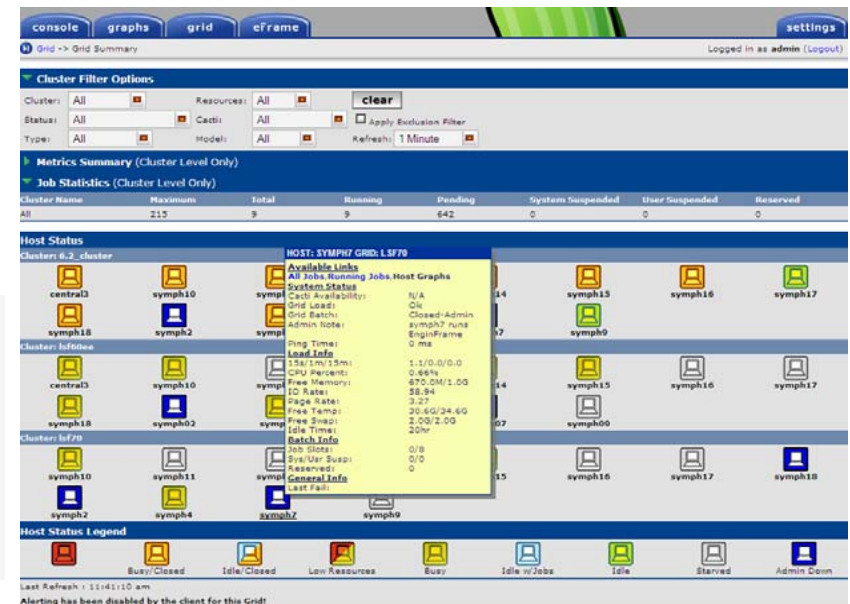




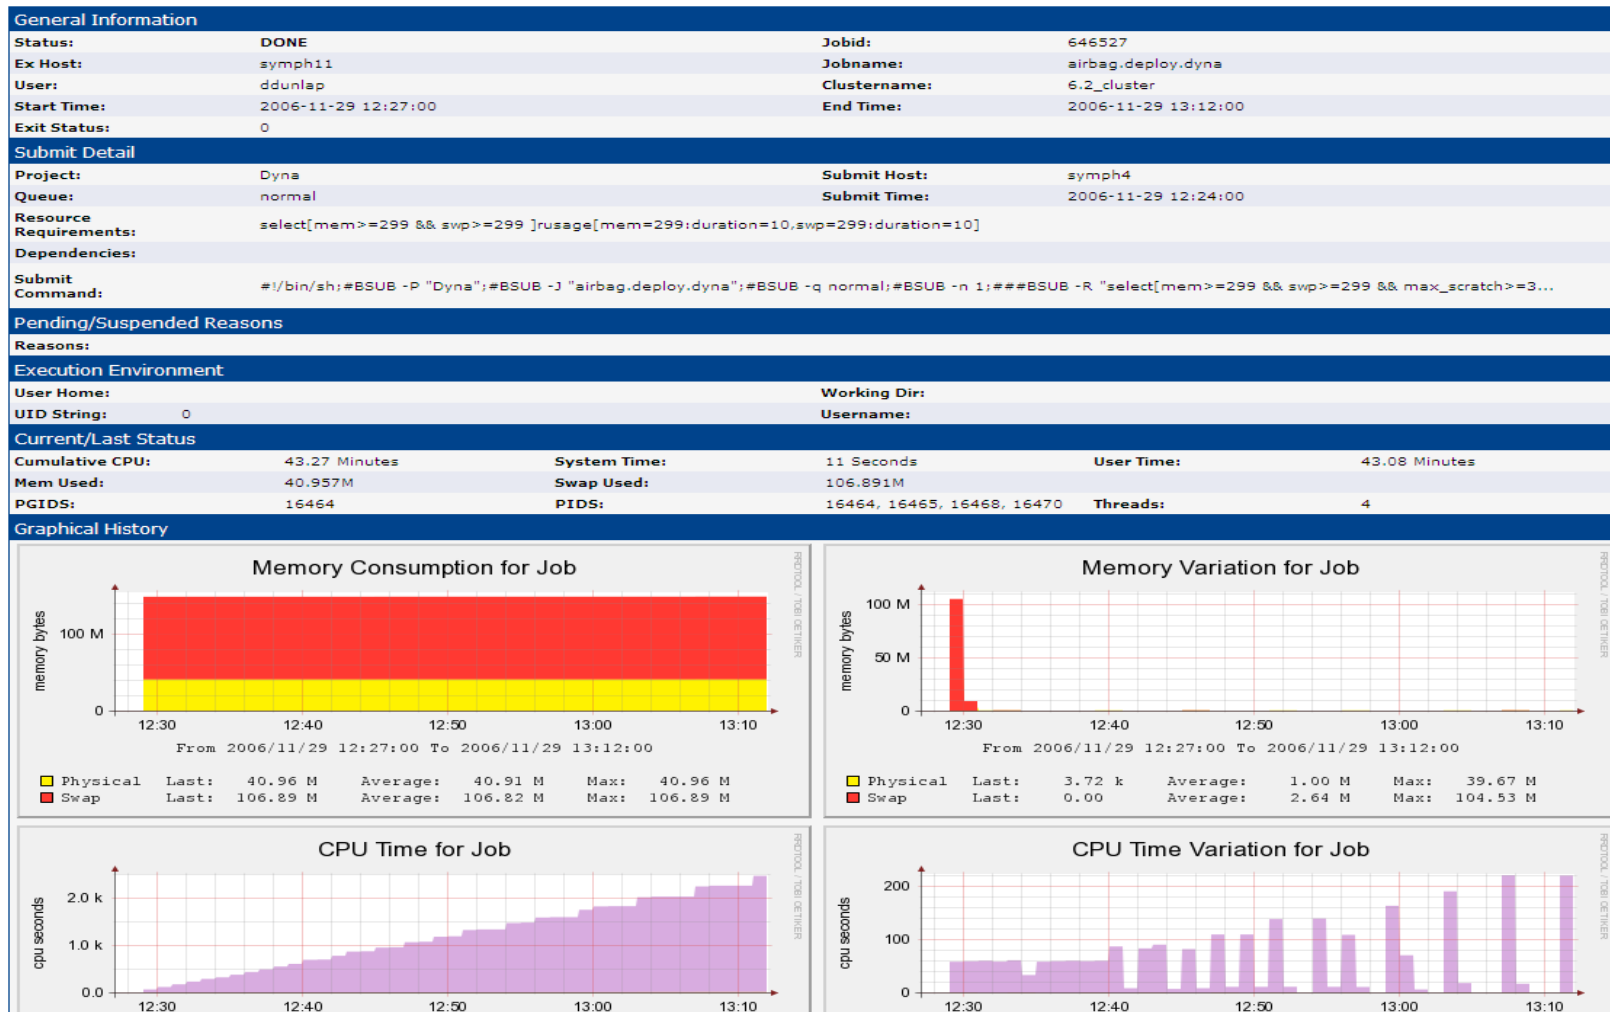
Platform Computing Simplifies

- Why we use Platform products: simplicity and flexibility
 - User base: hundreds of heterogeneous users
 - Mixed modes: from EP to MPI
 - Flexibility: open API
 - Monitoring: reporting and analytics
 - Ease of use: simple and clearly documented



- Tool for LSF Administrators and Operators
- Near Real-Time Monitoring for LSF and extendable to other physical devices and application software
- Drill-down to Individual job statistics
- It allows you to graphically see your clusters job history without the use of b-commands
- Based on Open Source Cacti
- *Platform RTM enables system administrators to make timely decisions for proactively managing compute resources against business demand ...*





- Simplify HPC cluster

HPC Management Software – for Single Linux Cluster

Platform Open Cluster Stack (OCS)

Intel Software Tools Kit
MPI & HPX Kit
OFED tools Kit
Intel Cluster Ready Kit

Platform Lava
PVFS2

Nagios, Cacti NTOP
Core Kusu Management
Tools

Develop

Schedule & Run

Manage

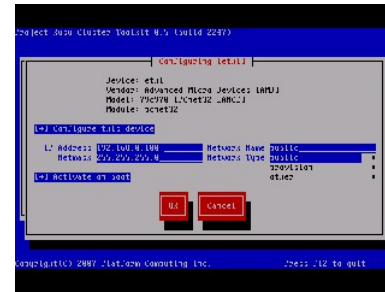
Red Hat HPC Solution = RHEL + Platform OCS 5




Boot from CD




Choose Language



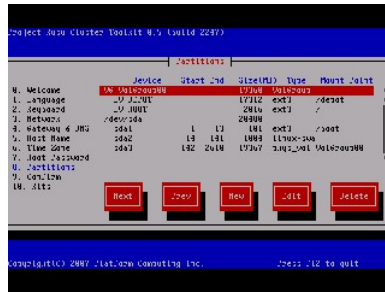
Configure Networks



DNS, & Gateway



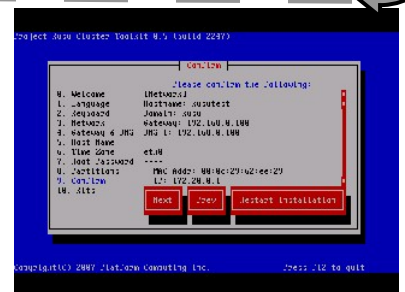
Root Password




Partitioning & LVM



Adding Kits



Install Summary



Installing Packages



Installer Node Boot



Creating an Installer Node

Installation complete.
Installer node is ready
to use.

**Operating Systems:**

RHEL 4, 5*
 Fedora 6
 Fedora 7
 CentOS 5

Message Passing Libraries:

MPICH 1,2
 MVAPICH 1
 OpenMPI
 Intel MPI*

Workload Management:

Platform Lava
 Platform LSF*

Resource Management:

Platform EGO*

Application Portals:

NICE EngineFrame*

Certification Tools:

Intel Cluster Checker+

Applications & Benchmarks:

Atlas
 blacs
 bonnie++
 fftw

hdf5
 iozone
 iperf
 linpack
 modules
 netcdf

scalapack
 stream

Networking Hardware:

Gigabit Ethernet
 Cisco Infiniband
 Qlogic Infiniband
 Mellanox Infiniband

Cluster Provisioning Methods:

Package
 Diskless
 Imaged

Repository Management:

Multiple repositories
 Repository Snapshots
 Full Dependency Checking
 Package management

Cluster Update:

Update from RHN
 Update from Yum repository
 Update individual rpms

Configuration Management:

Node Templates/Classes
 Cluster File Synronization
 Driver Management
 Cluster Network Configuration

Cluster Administration

Cacti Configured to monitor
 Memory on a node
 CPU Load on a node
 Logged In users
 Processes on a node

logwatch configured to monitor activity on the cluster

Pdsh for administering nodes on the cluster

Nagios configured to manage & Notify

Nodes in the cluster
 Web server
 Load
 logged in users
 DNS Server status
 MySQL Server status
 NTP daemon status
 Node 'PING' status
 Root Partition disk
 SSH Daemon
 SWAP Space
 Total Processes on node
 Email Notifications to Administrator

Plone CMS containing:

Manual pages for OCS
 Documentation for
 Platform OCS 5
 List of installed Kits on the cluster

*Optional Commercial products that must be purchased separately

+Available only from Intel Cluster Ready Partners



- Workload based dynamic power management
 - Turn on/off machine
 - Avoid hot spot
 - Slow down CPU speed when MPI task is waiting.
- Workload driven dynamic provisioning
 - Switch machines between Windows and Linux automatically based on workload
- “Cloud” Toolkit
 - User web portal
 - Dynamic provisioning
 - Workload scheduling
 - Accounting and billing



24x7 Support across the globe

“Platform has been proactive, involved and very, very friendly in providing support.”

Henry Neeman
*Director, Oklahoma University
Supercomputing Centre*

“Platform’s standard of support has been excellent.”

Tim Cutts
*Platform LSF Administrator
Sanger Institute*



- Platform's Value proposition
 - Total solution for HPC management
 - 16 years of HPC experiences and broader customer base
 - Top quality technical support and services

Platform™

Powering High Performance

www.platform.com

info@platform.com

1-877-528-3676 (1-87-PLATFORM)