Platform™
The Power of Sharing

Cloud Computing for HPC

William Lu
Platform Computing

Platform

Trends

- Increasing demand for compute power in simulation, analysis, and research
  - More Data, More Compute, More Analysis
- Organizations are searching for an integrated management stack to further improve user service level and simplify HPC operations

Platform

Plans to Deploy Private Cloud in 2009

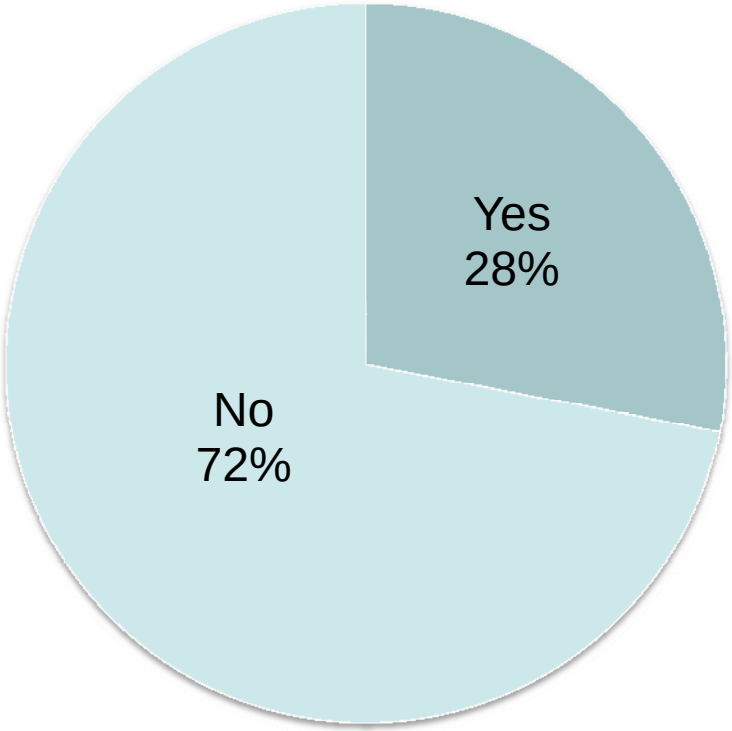

Source: International Supercomputing Conference 2009 attendee survey

Motivations for Private Cloud

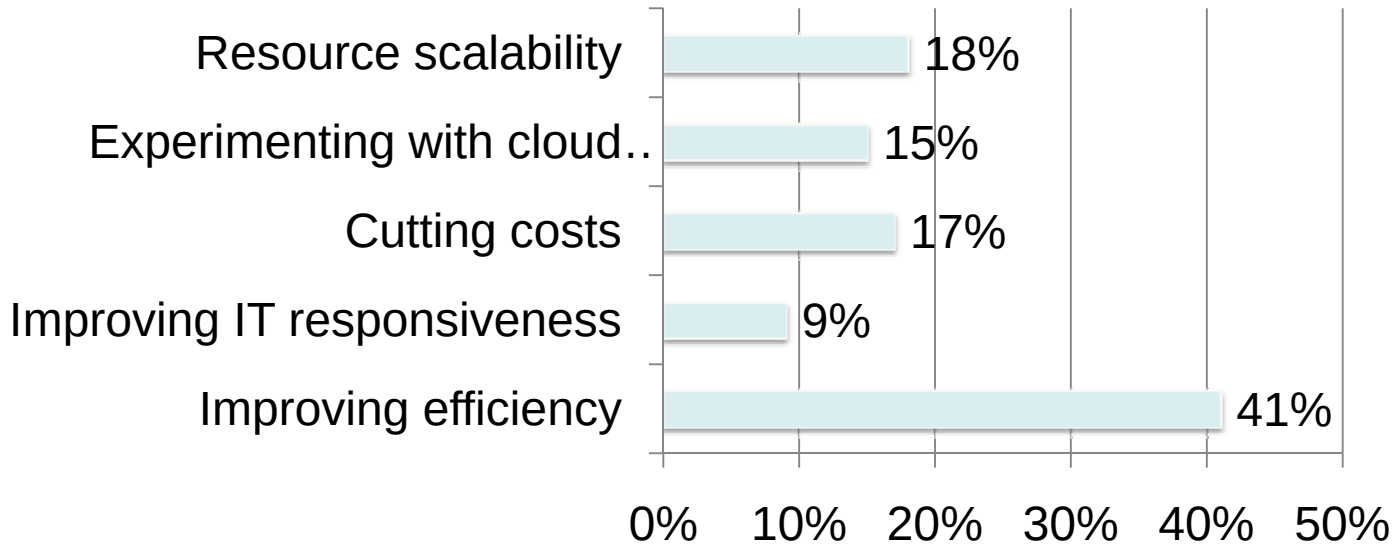

Source: International Supercomputing Conference 2009 attendee survey

Barriers to Private Cloud

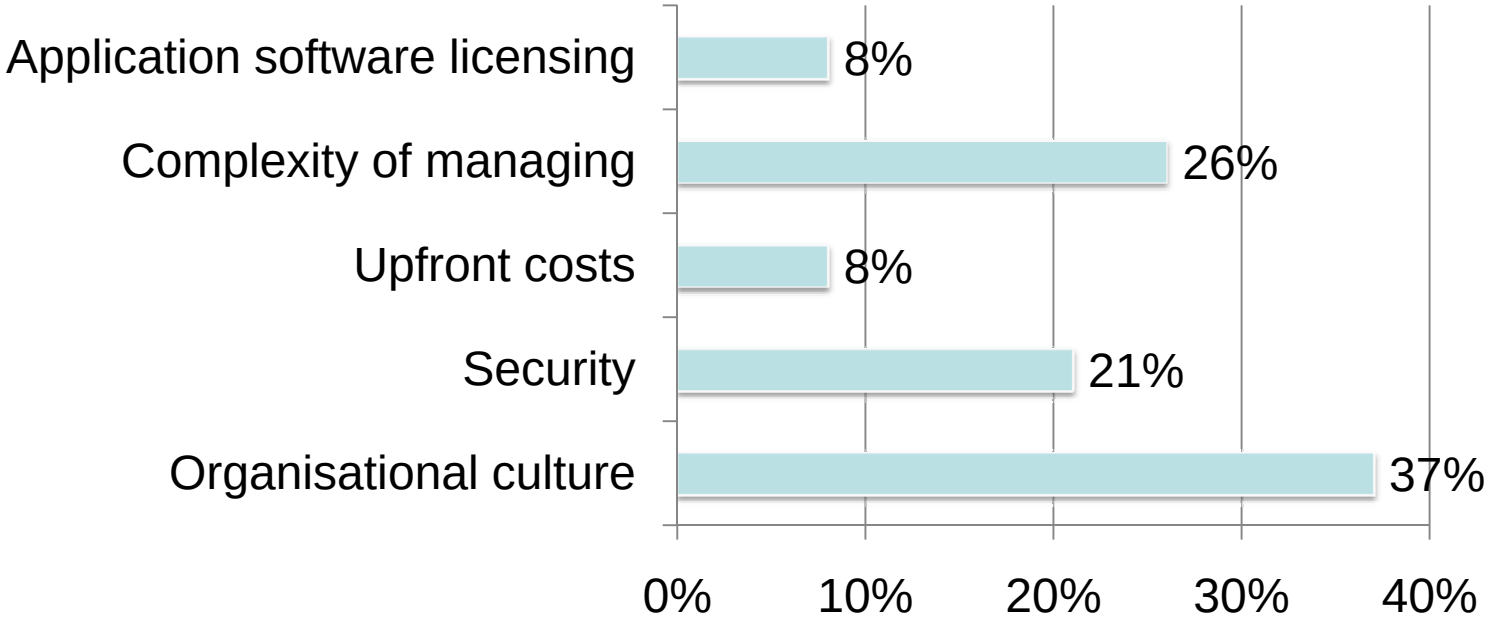

Source: International Supercomputing Conference 2009 attendee survey

Platform

Challenges in HPC grid/clusters

- Job interference
  - Memory leak
  - Run time environment (OS+tools)
- Security
- “One fits all” scheduler
  - Fair share with history vs. current resource allocation

Reality Behind Resource Allocation

- Static (advance) Allocation vs. Real Workload

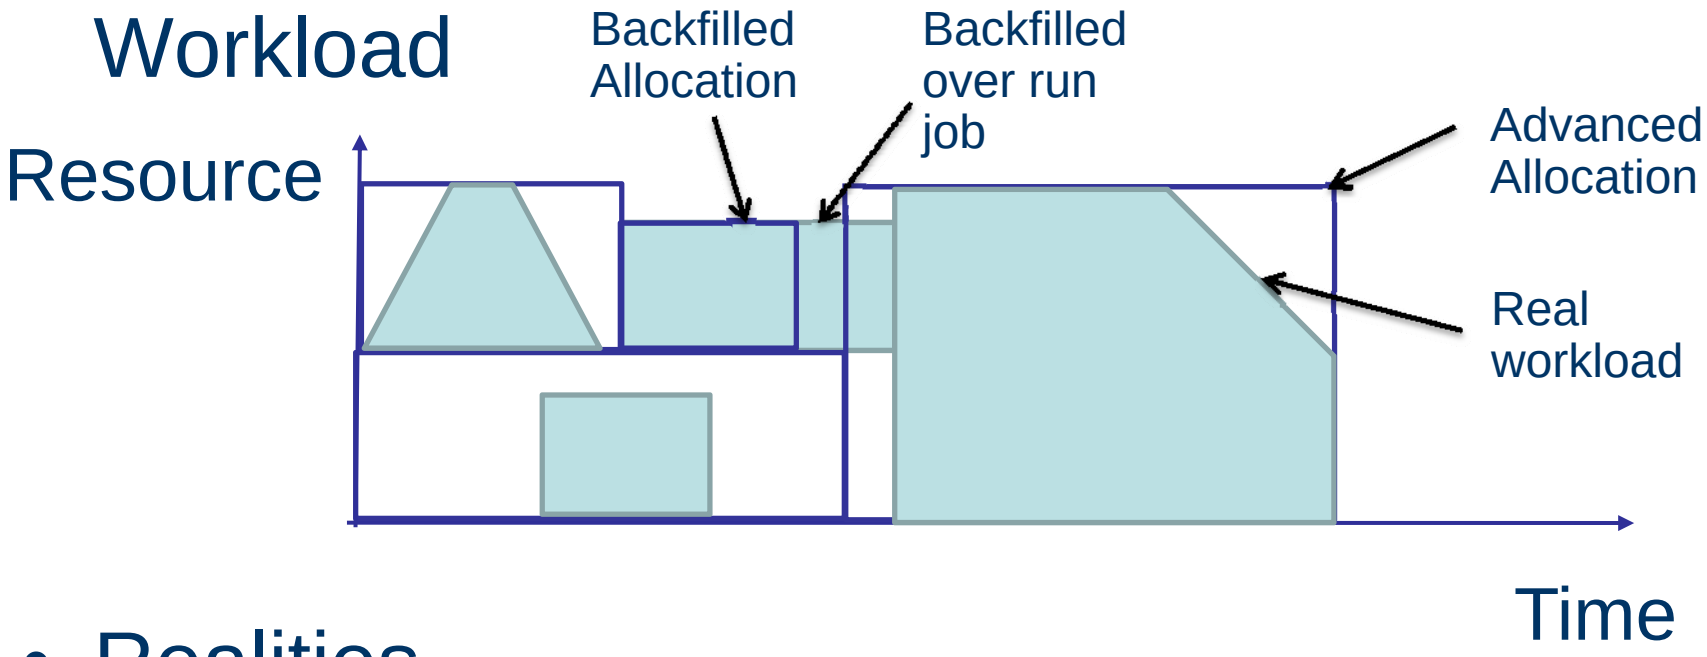

- Realities
  - Lacking of visibility of real workload
  - This delivers SLA/QoS
  - But… real utilization is low

Elastic vs. Traditional Scheduling

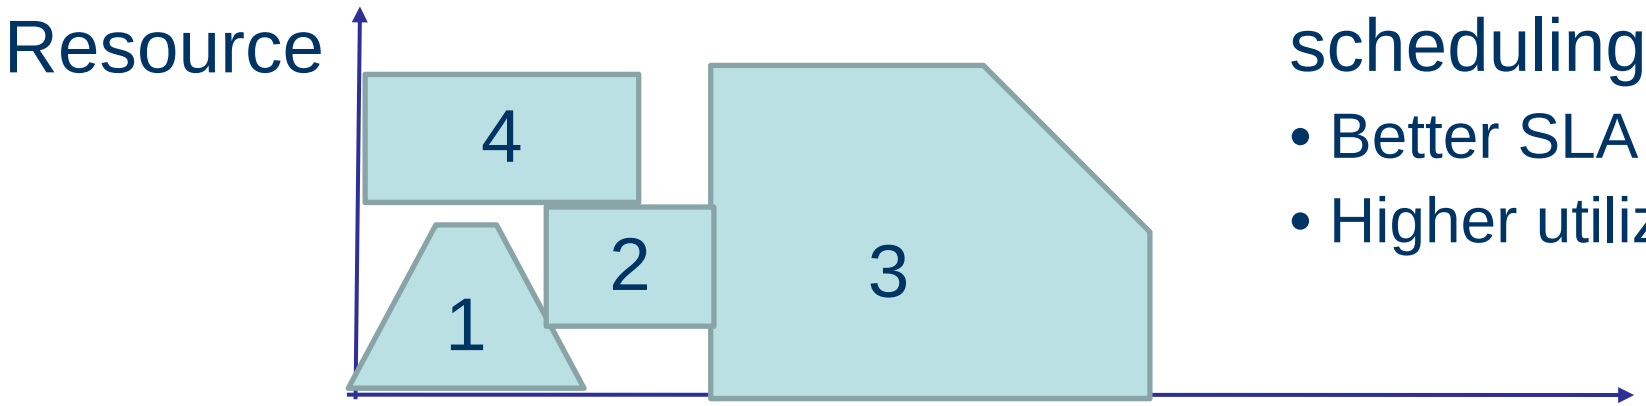

Elastic scheduling

- Better SLA
- Higher utilization

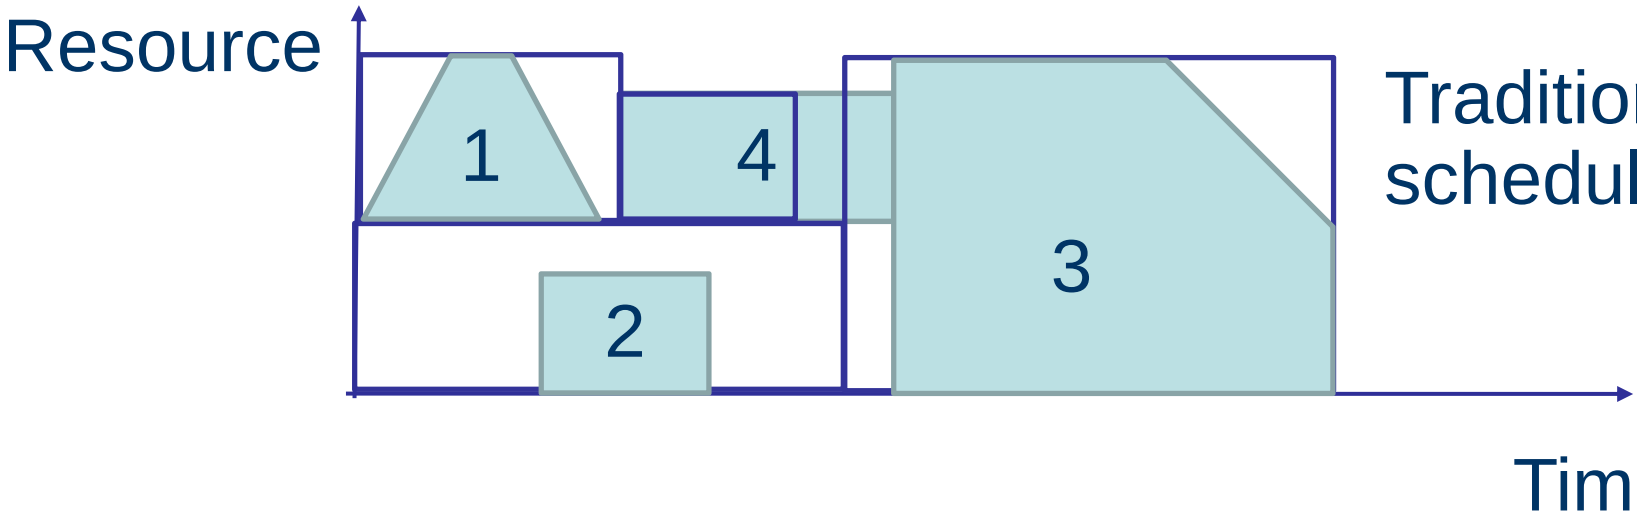

Traditional scheduling

Resource Scavenging

- Scavenging resources in desktop, labs, external data centers, public cloud…

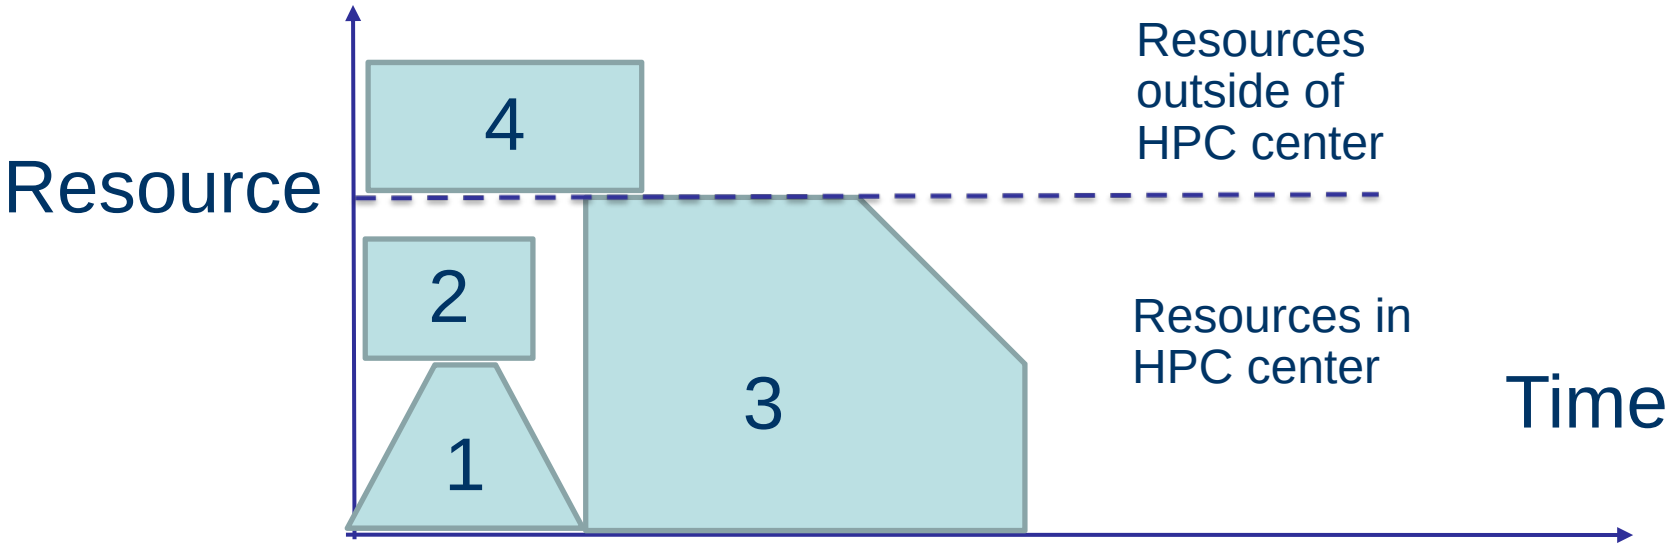

Different job environment

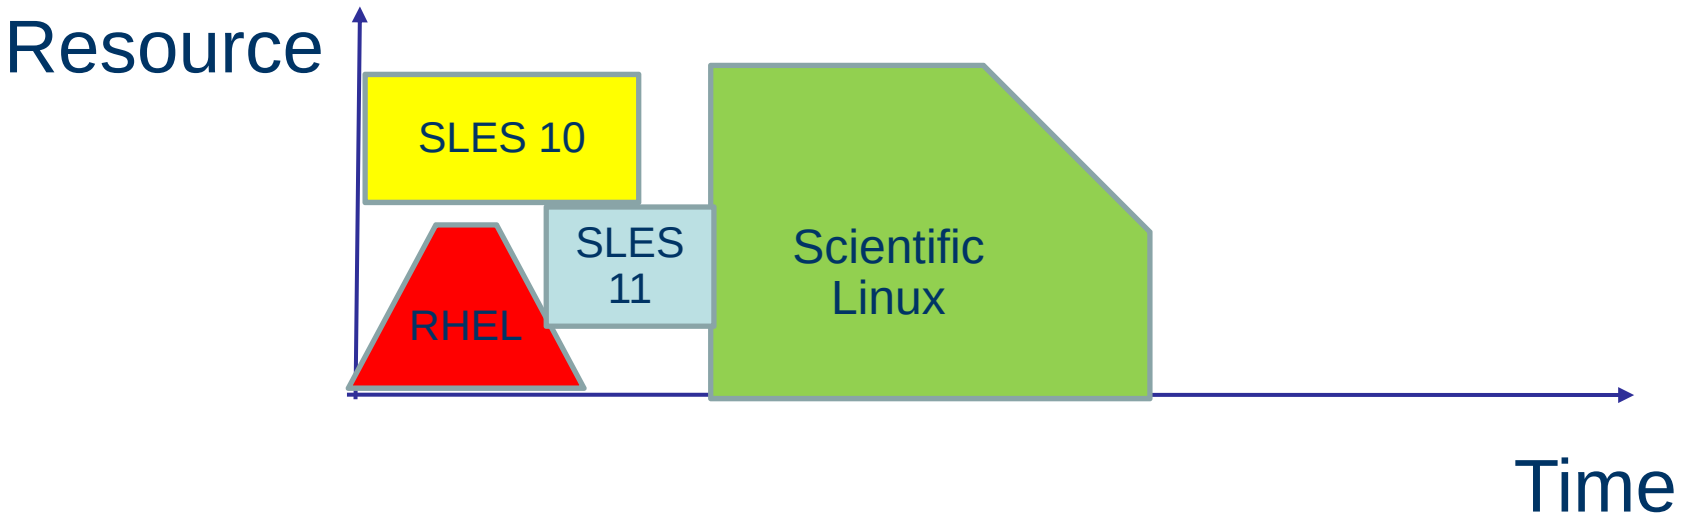

- Job runtime environment needs to be dynamically scheduled together with the job

Moving from Cluster/Grid to Cloud

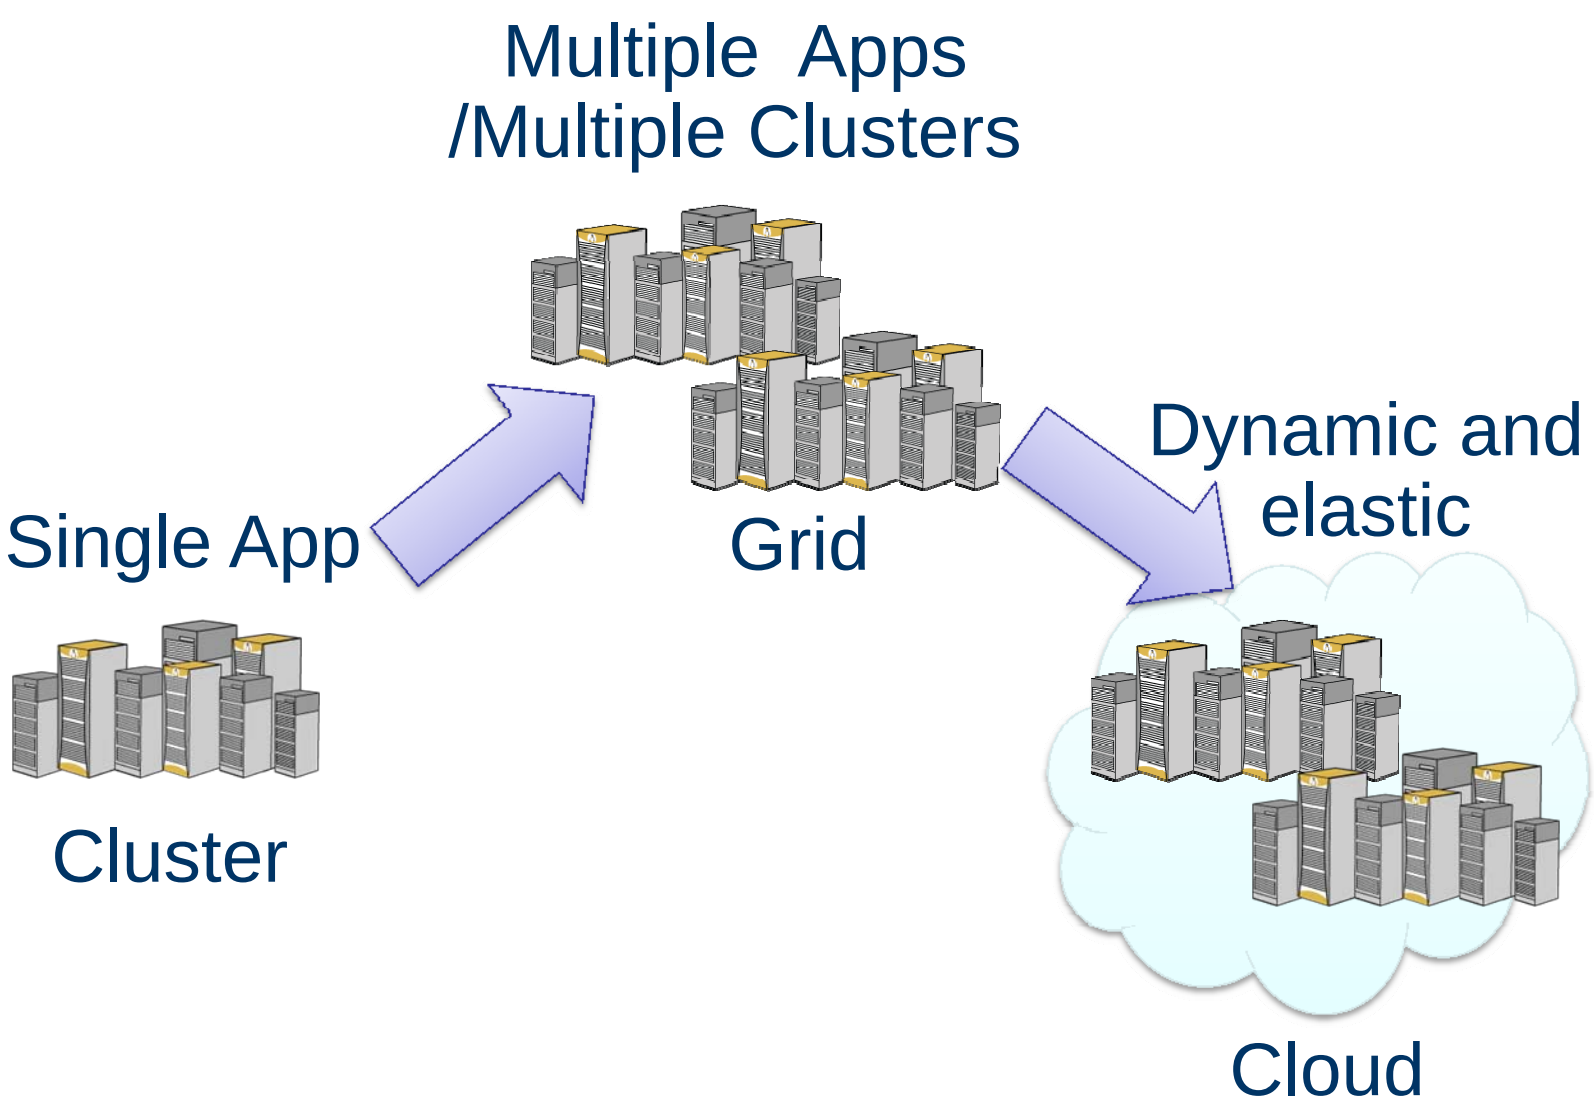

Management of Private Cloud for HPC

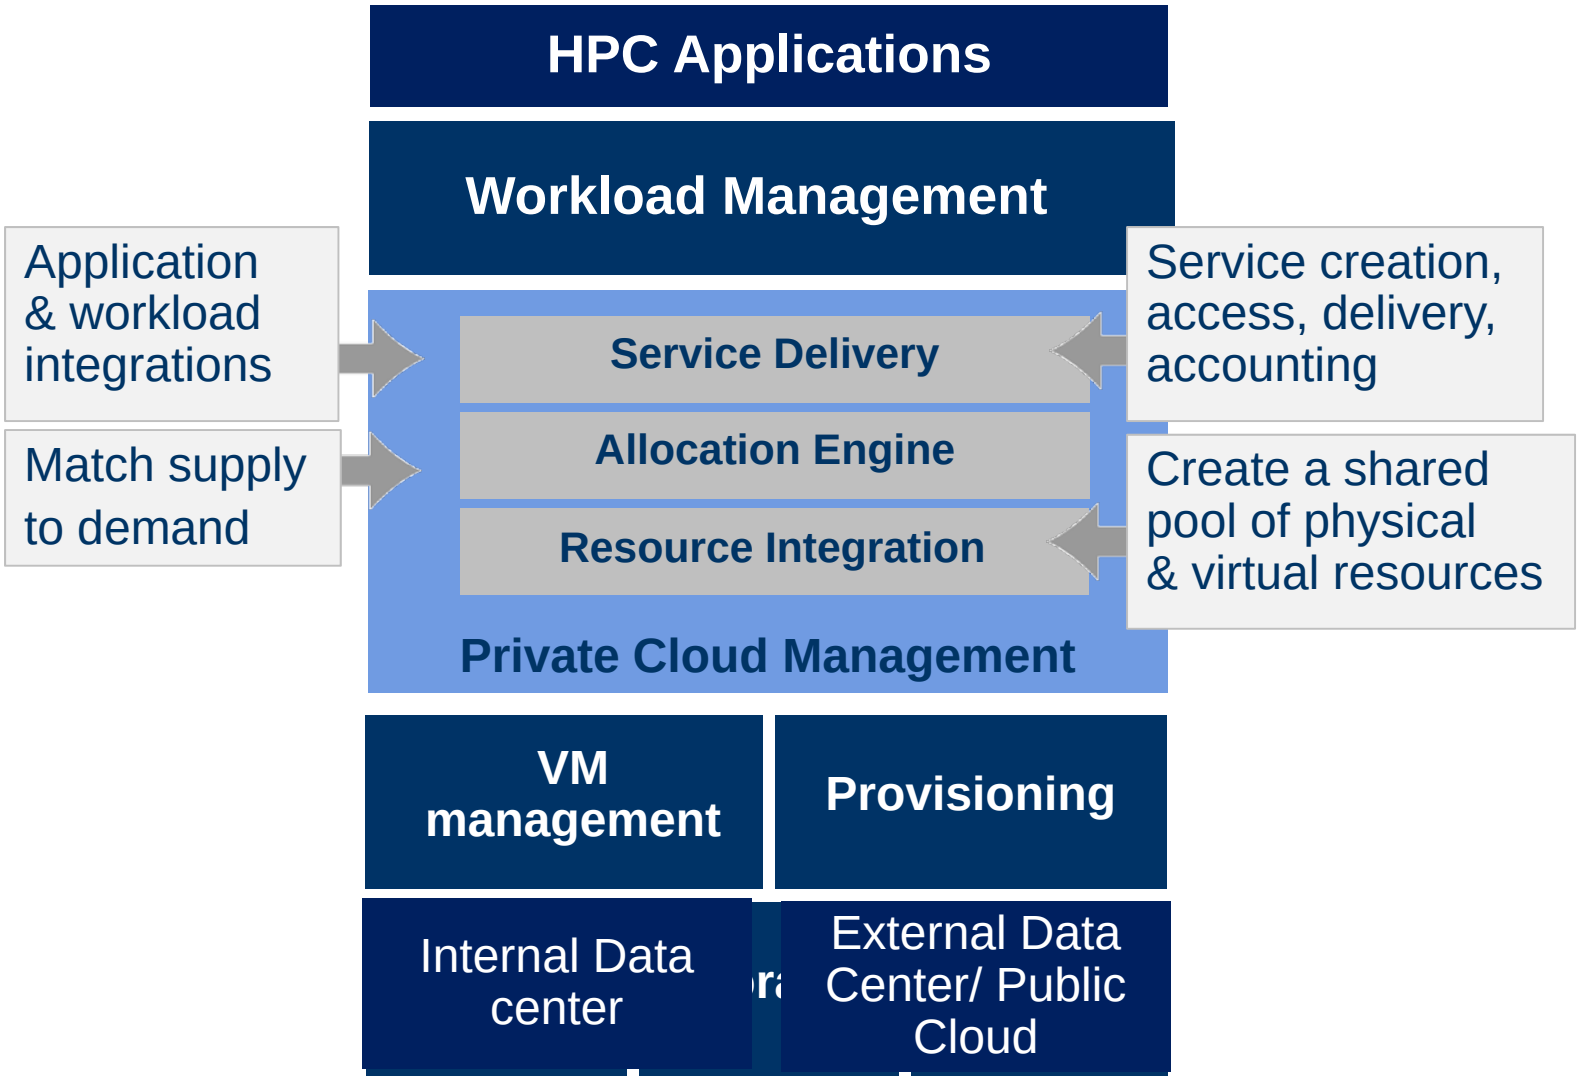

Software Stack from Platform Computing

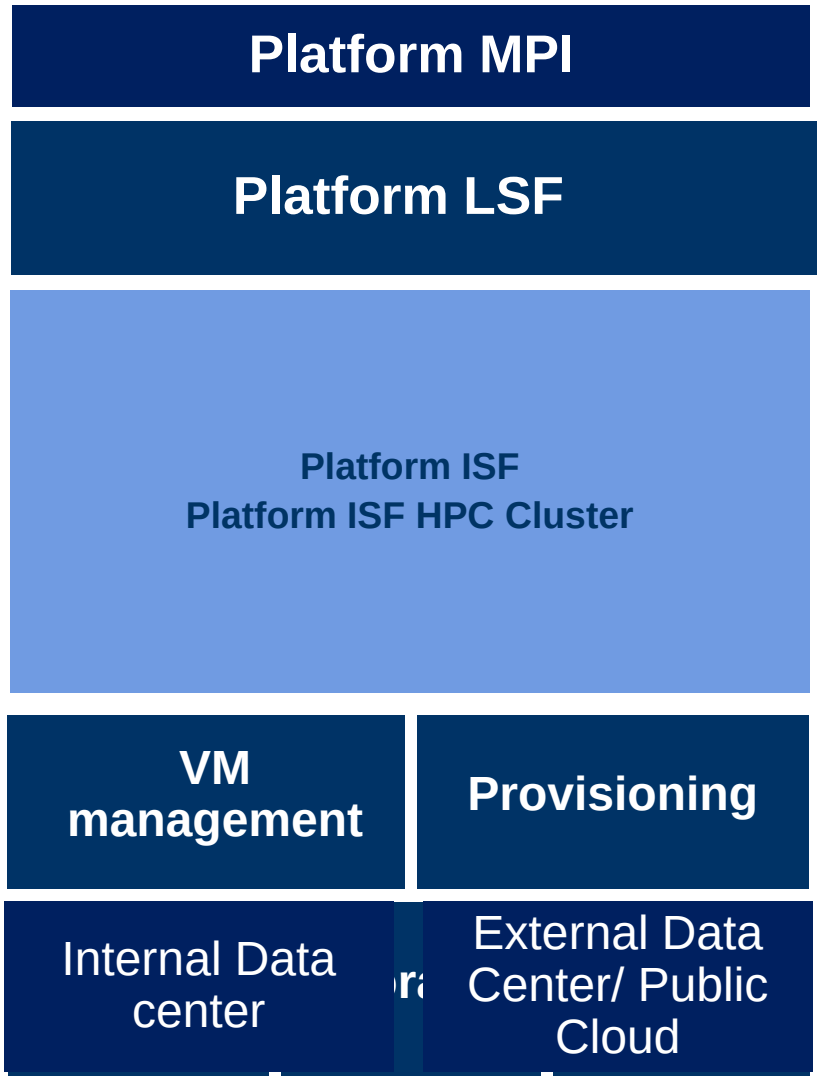

Large Research Institute

Profile

HPC center with thousands of servers used by multiple user groups for data analysis

Problems

- Application environment set up is too long
- Job fails due to application environment conflicts on the same server
- Scheduling policy conflicts across multiple user groups causing user dissatisfaction

Solution

- Elastic and autonomous Platform LSF clusters scheduled by Platform ISF for each user group
- When workload changes, Platform ISF dynamically add and remove resources from individual clusters
- VM is used as job container to reduce job interference

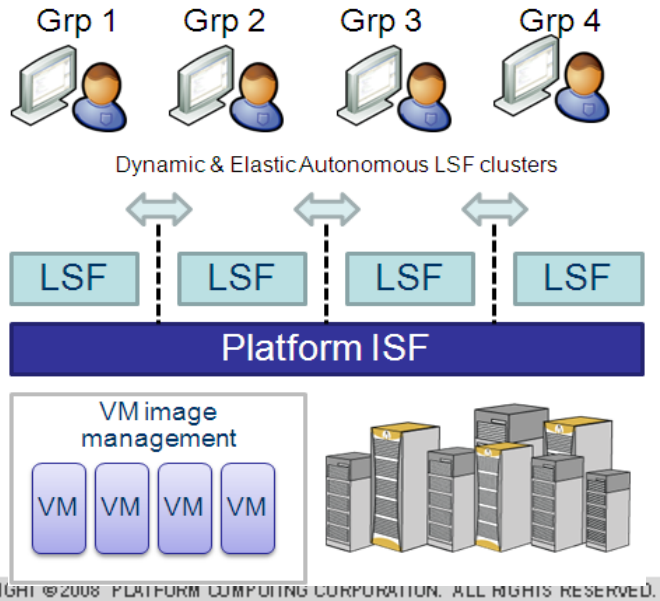

	Before	After
User support incidence	15/day	2/day
Job failure due to app environment conflict	5/day	~0
App env set up time	1 weeks	1 day

Reservation & On-Demand Allocations

On-Demand:

- First-come, first server capacity
- Pay-per-use for non-critical or burst workloads

Guaranteed Reservation:

- Reserve resources for critical workloads
- Optimization utilization, guaranteed capacity

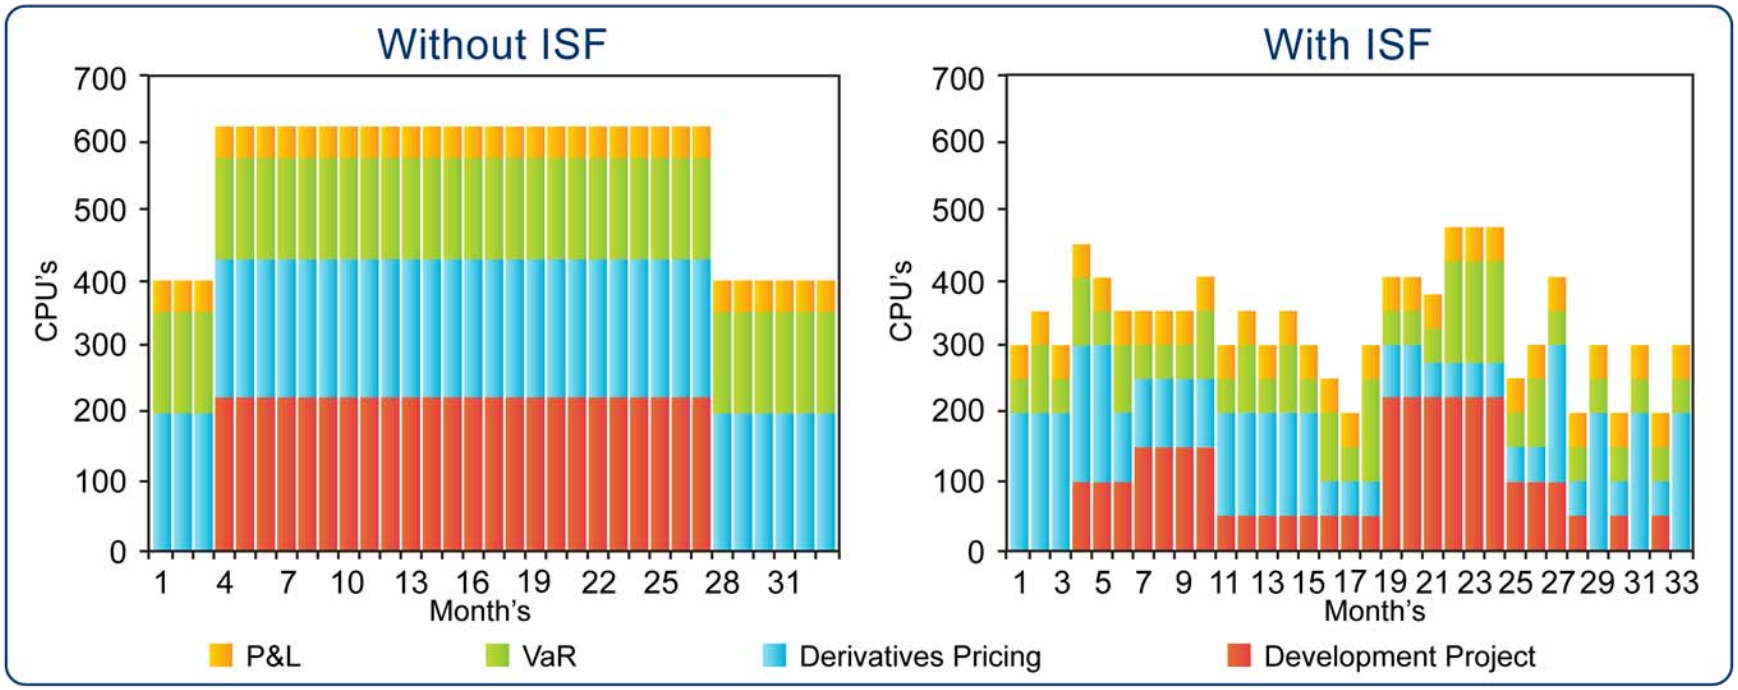

Resource-Aware Allocation Policies

Packing	▪ Pack workload on fewest number of physical servers ▪ Maximizes usable capacity, reduces fragmentations, reduce energy consumption	
Striping	▪ Spread workload across as many physical servers as possible ▪ Reduce impact of host failures, higher application performance	
Load-Aware	▪ Allocate physical servers with lowest load to new workloads ▪ Higher application performance	
HA-Aware	▪ Allocate HA-enabled resources to critical workloads ▪ Match availability levels to service requirements and costs	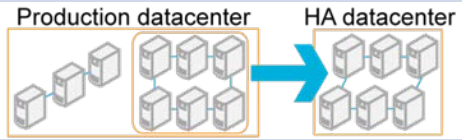
Energy-Aware	▪ Place workload according to energy indices and datacenter hot spots ▪ Reduce energy consumption	
Affinity-Aware	▪ Place workload close to critical resources such as storage ▪ Higher application performance	
Server Model- Aware	▪ Allocate resource to workload according to model types ▪ Maximize utilization of higher performing & more expensive resources	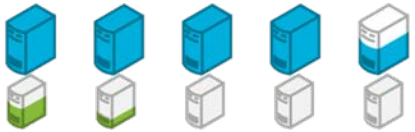
Topology-Aware	▪ Allocate resources on the same interconnect to the same application ▪ Improve application performance	

Reporting and Accounting

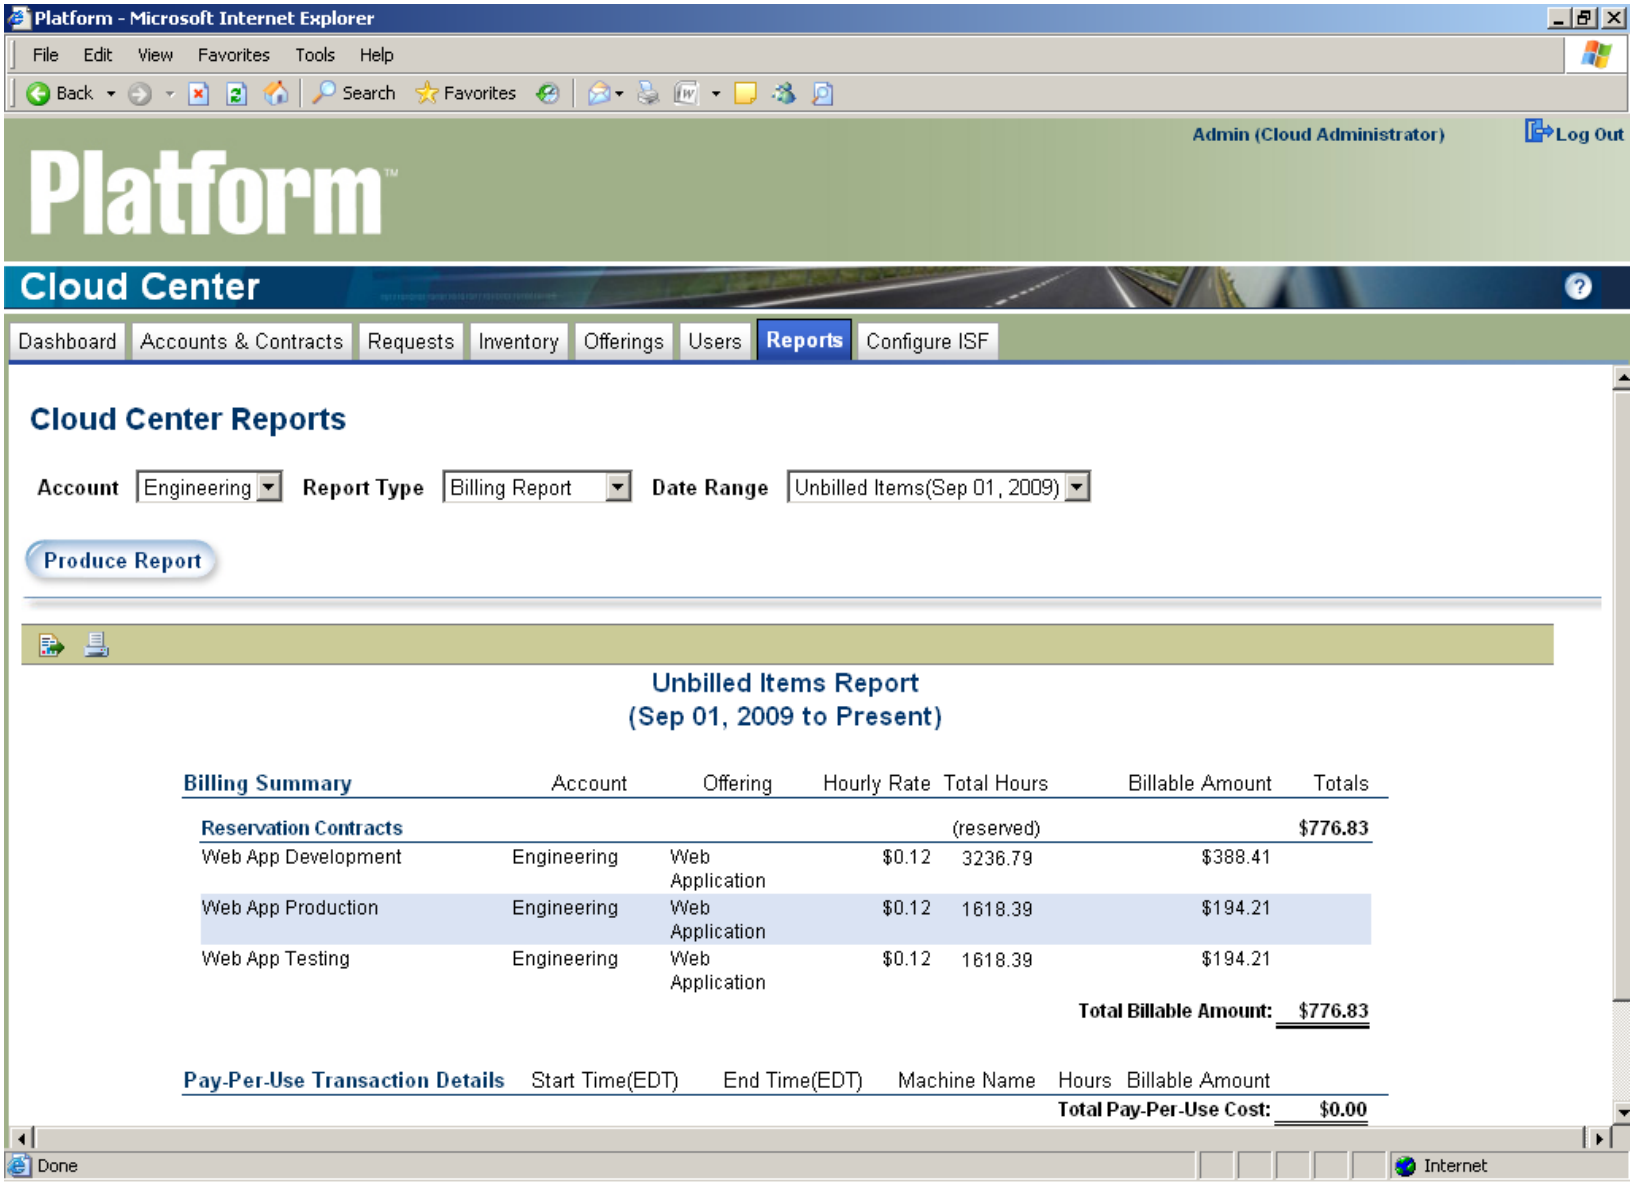

User Access Point – LSF Web Portal

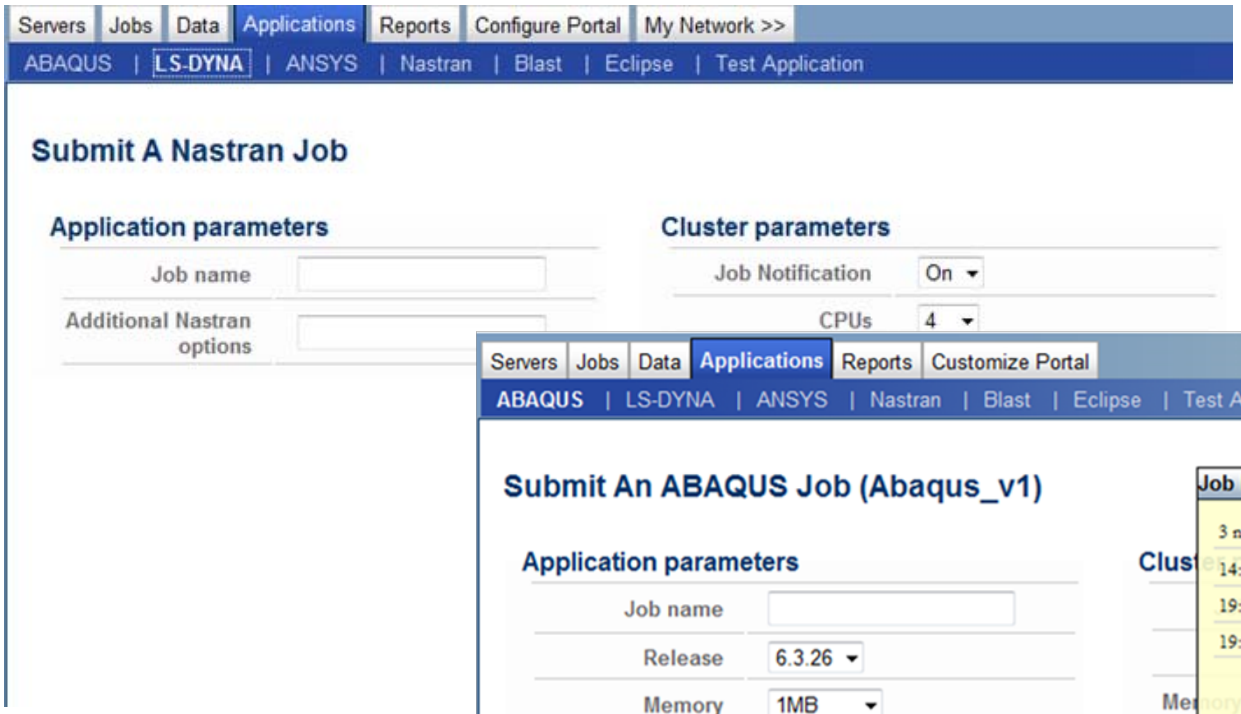

App Centric Interface

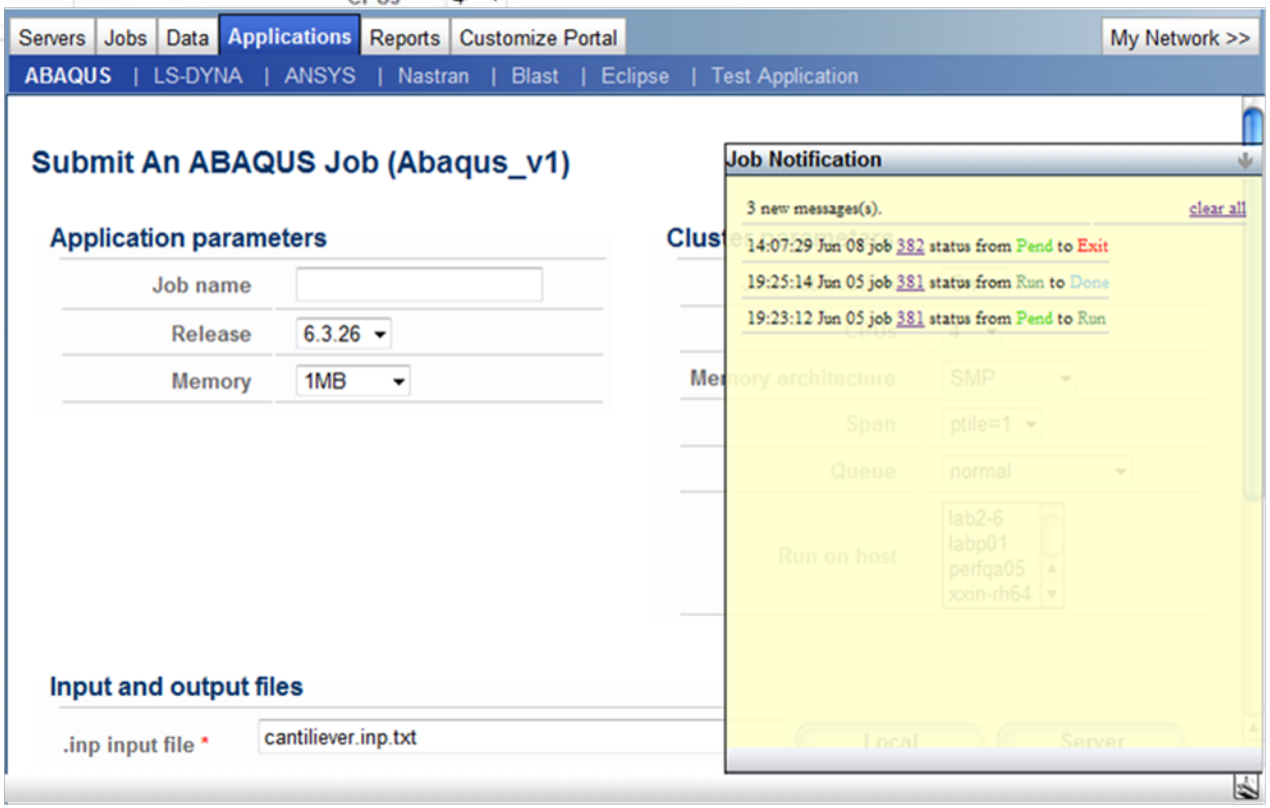

Job notification

User Access Point – LSF Web Portal

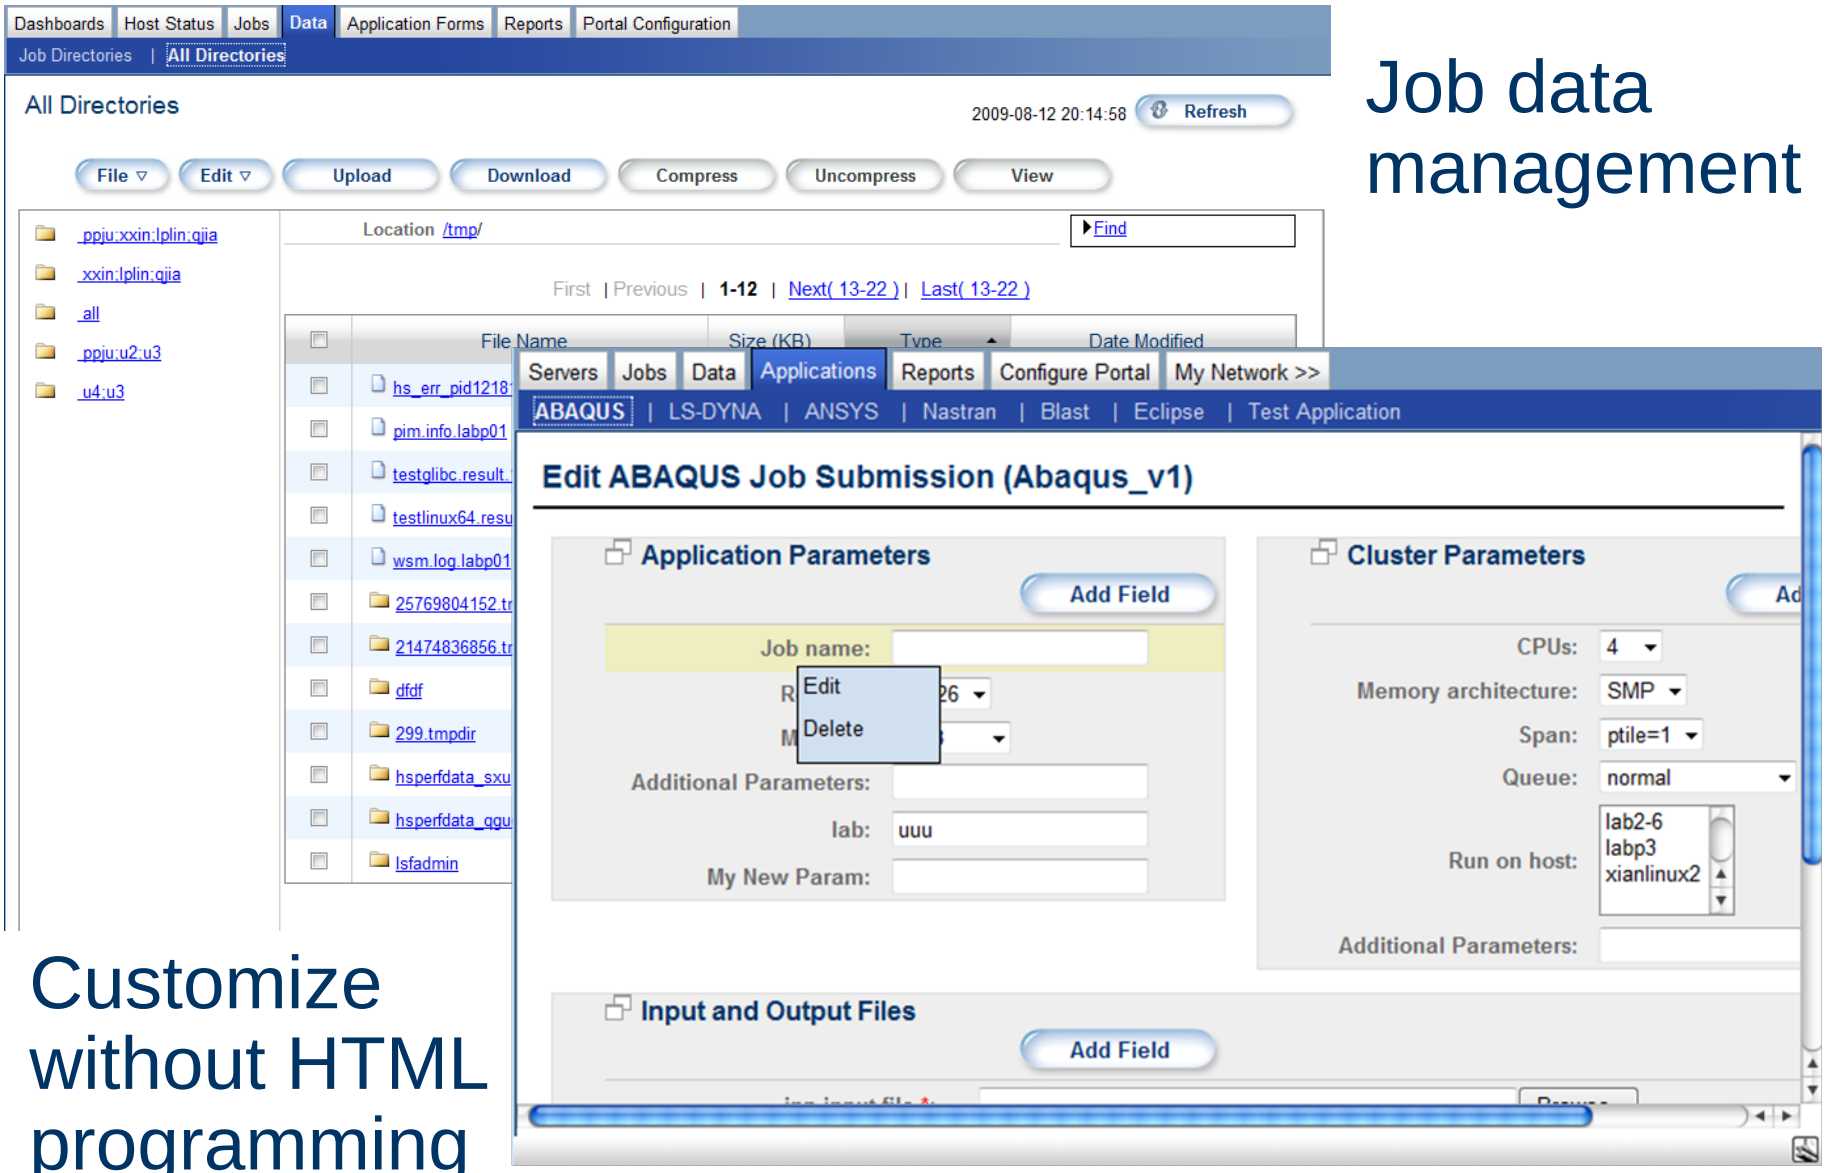

Job data management

Customize without HTML programming

Platform

Summary

- Cloud is happening due to its superior SLA
- Cloud is to add dynamics and elasticity to grid and cluster environment
- Platform ISF delivers elastic logical clusters in a shared HPC infrastructure. It transforms cluster and grid to cloud

Platform™

The Power of Sharing

Thank You

www.platform.com