

HPC in a Cloud

June 2009



- **Bridge the gap between HPC usage, users and its potential**
- **Provide added expertise to OEMs/users to help with understanding, operating and optimizing HPC systems**
- **Guide users on how to improve efficiencies in research, education, innovation and product manufacturing**
- **Provide application designers with the tools to enable parallel computing**
- **Strengthen the qualification and integration of HPC system products**
- **Enable new users to leverage from the experience of HPC experts**
- **Review future HPC technologies and facilitate discussions on their development**

For more info please visit: <http://www.hpcadvisorycouncil.com/>

HPC Advisory Council Activities



Home | Blog | Council Members | Cluster Center | Network of Experts | Technical Content | Contact

Mellanox Cluster Center
Apply for cluster access

HPC Advisory Council
Mellanox is dedicated to building a distinguished HPC alliance by working closely with our chosen partners and customers to ensure the best total solution is available to end-customers. The HPC Advisory Council includes best-in-class original equipment manufacturers (OEMs), strategic technology suppliers, independent software vendors (ISVs) and selected end-users across the entire HPC market segments.

Performance on Overdrive

The HPC Advisory Council is also a community effort support center for HPC end-users, providing the following capabilities:


- Mellanox Cluster Center - the center provides a unique ability to access the latest Mellanox and the HPC Advisory Council member technology, even before it reaches the public availability. It provides the Council members and any HPC end-user with a development, testing, benchmarking and optimization environment.
- HPC Advisory Council support group - provide a support center for consultations, questions / issues etc. for the HPC end-users.
- JOIN TODAY! To become an HPC Advisory Council member please refer to the HPC Advisory Council Application (PDF)
- READ THE HPC ADVISORY COUNCIL BLOG
- Current Member Roster

Network of Expertise

Best Practices

- Oil and Gas
- Automotive
- Bioscience
- Weather
- CFD
- Quantum Chemistry
- and more....

HPC Outreach and Education




Cluster Center
End-user applications benchmarking center



HPC Technology Demonstration

Over 80 members



40Gb/s InfiniBand Distributed Visualization over SCinet

40Gb/s InfiniBand SCinet Participants



HPC Systems as Application Systems



Common Practice:
HPC resources are acquired for specific purpose
They are typically dedicated for single type of work

Capacity limit

- The total capacity is limited by the size of the system

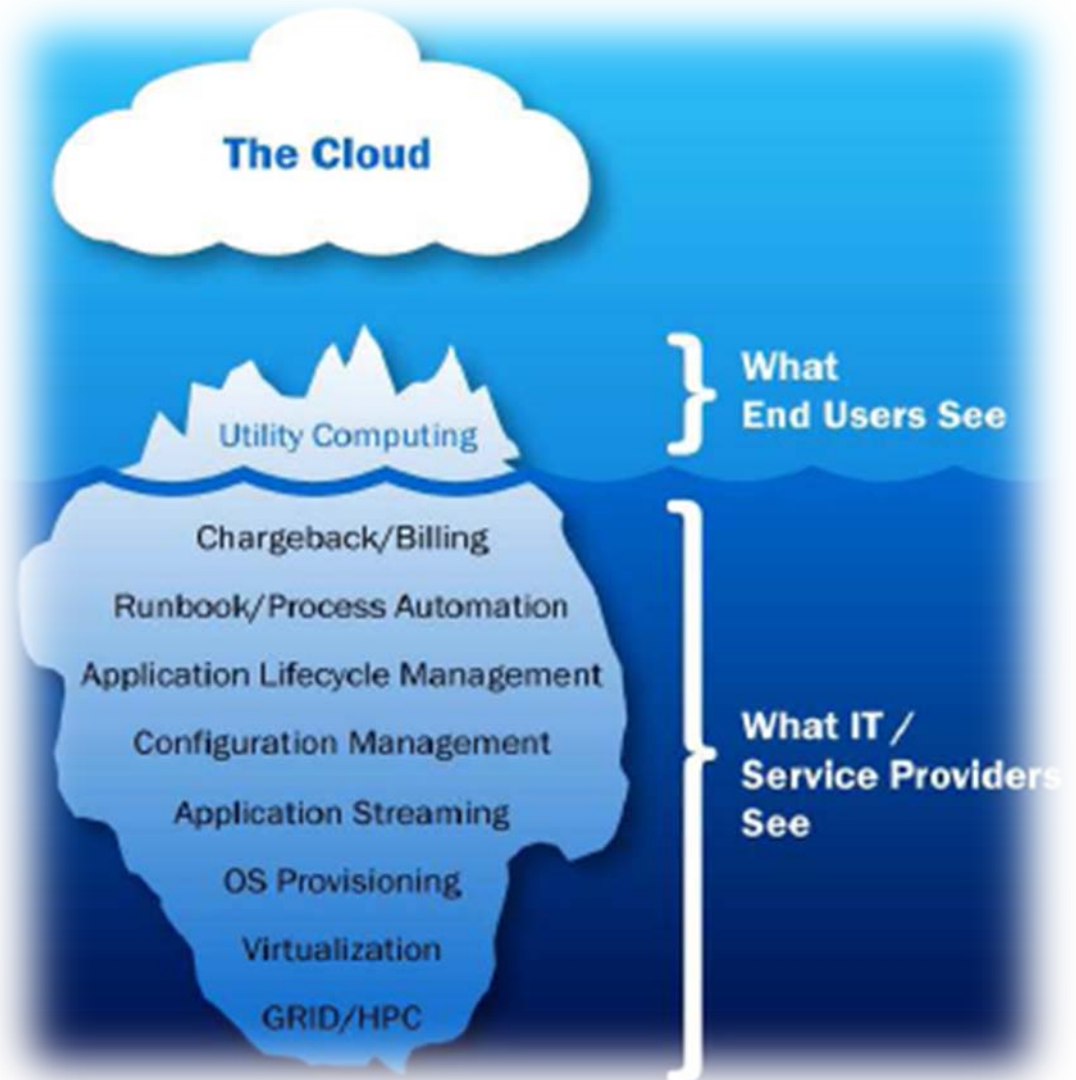
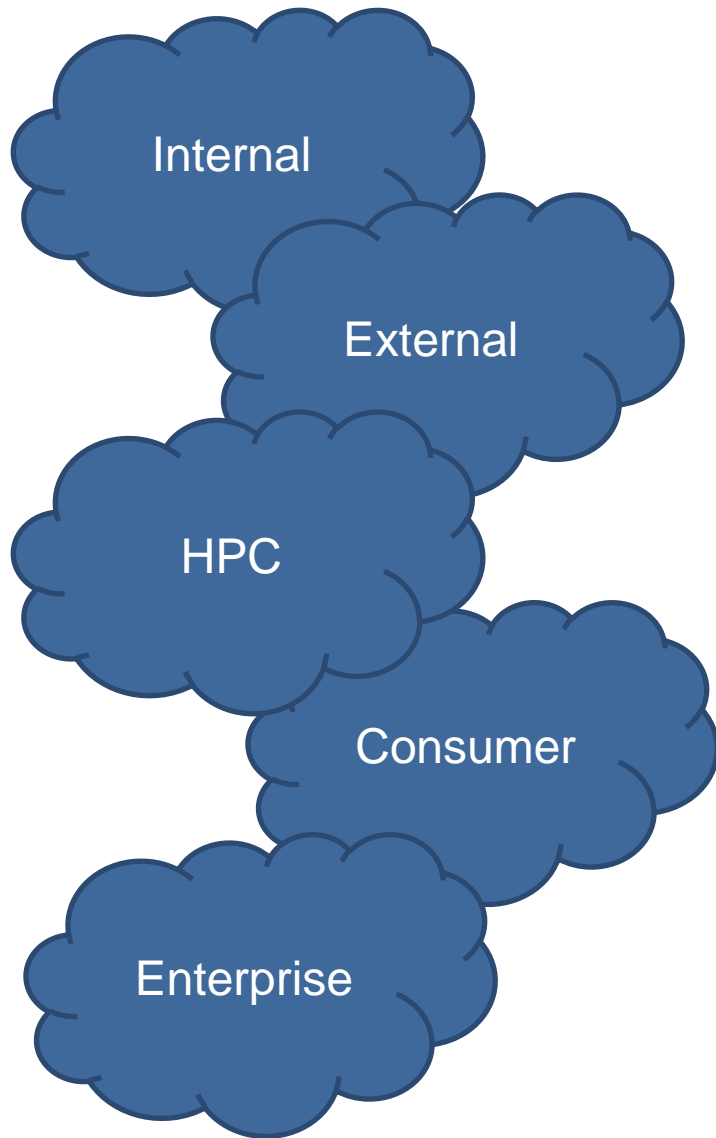
Utilization

- Provisioned for peak load
- Even if it is not fully utilized, it can't be repurposed for other applications

Quick Resource Provisioning

- Users compare internal HPC resource with external

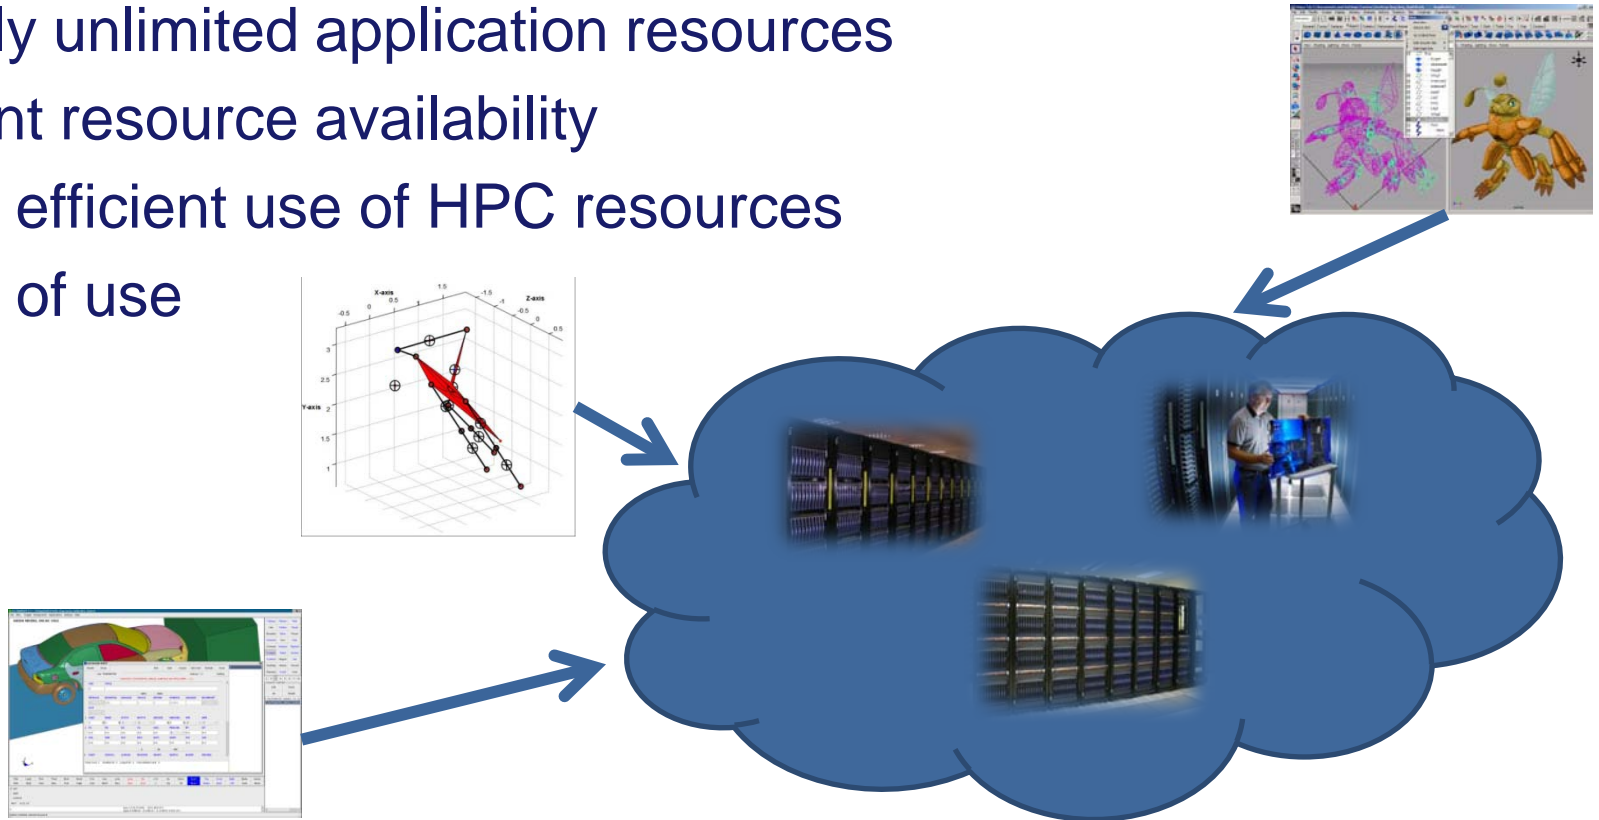
The Cloud Computing Concept



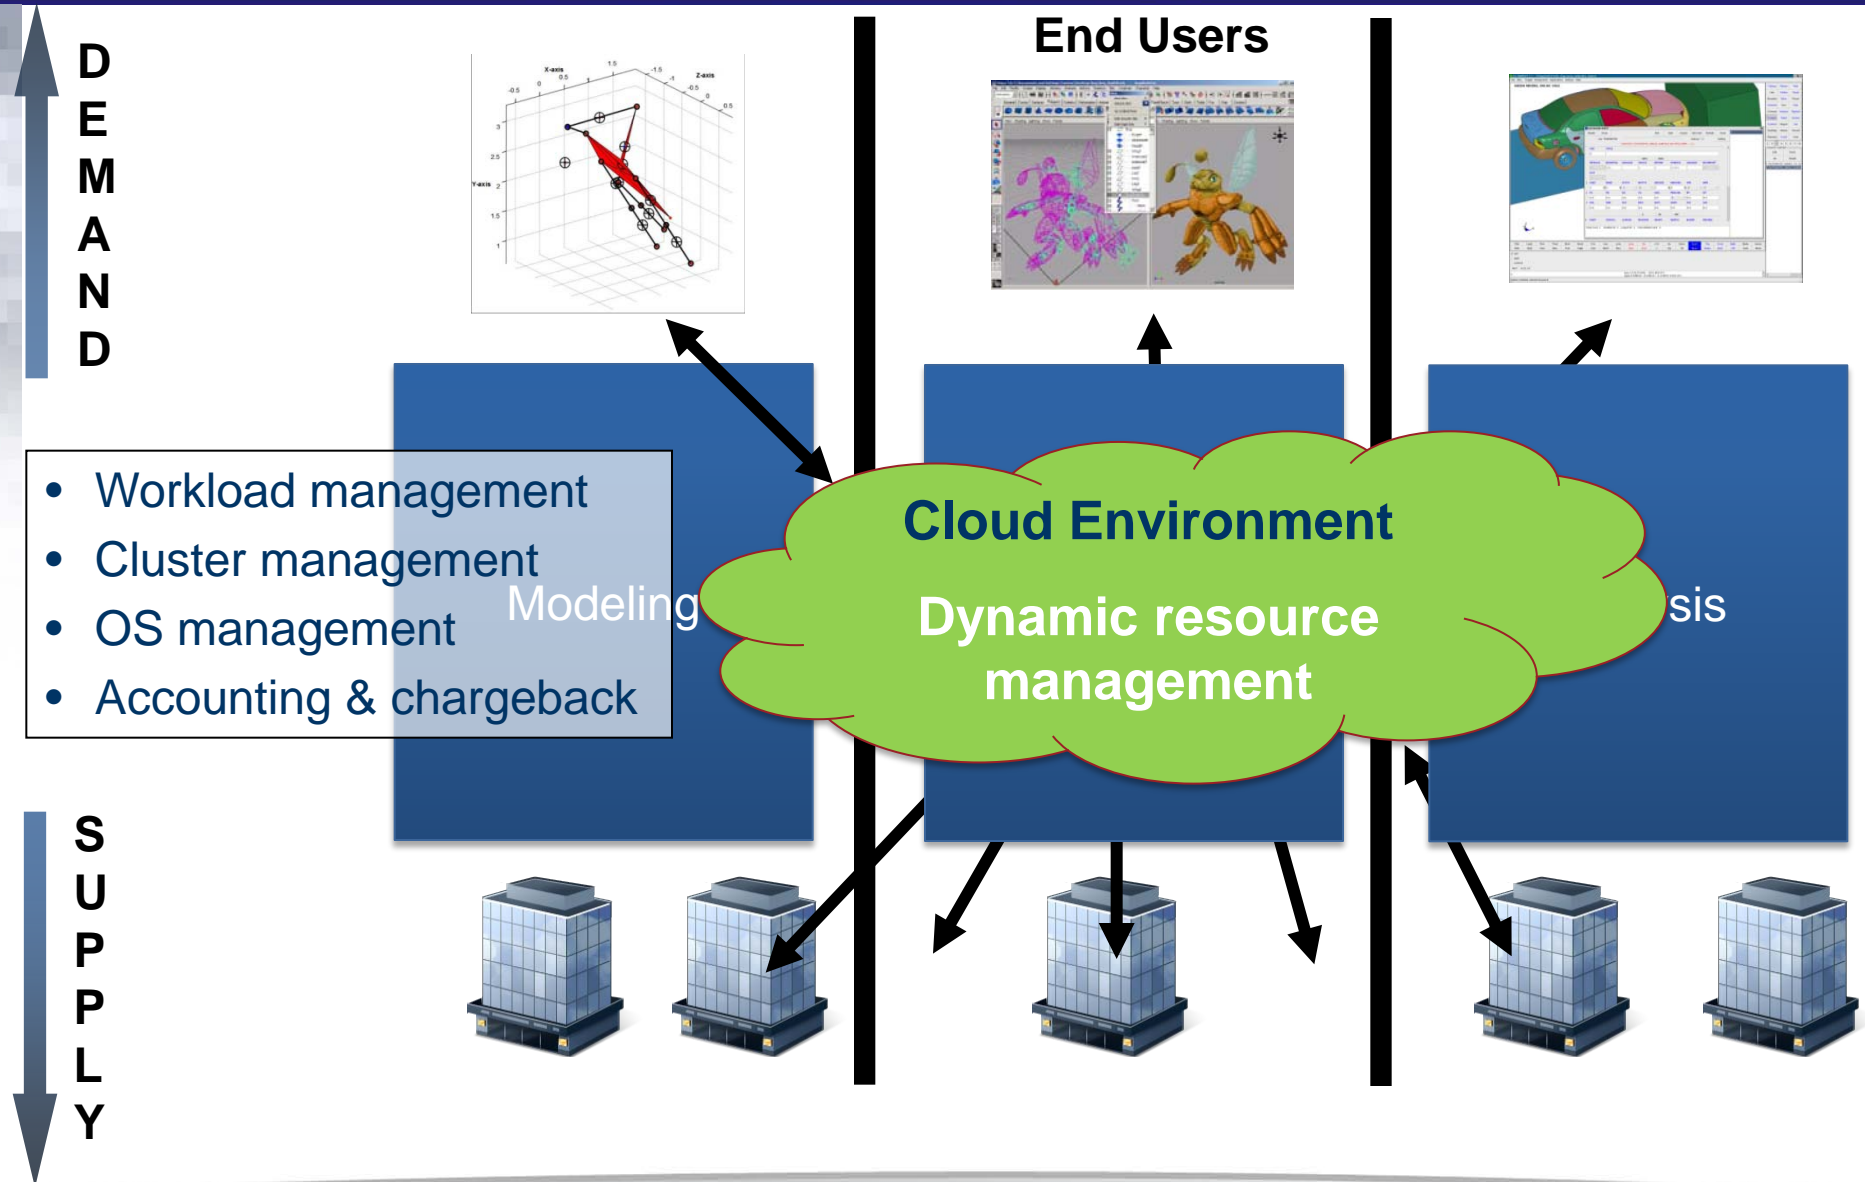
- Image Courtesy 451 Group

The Concept of HPC Cloud

- **Software and Hardware** as a Service
- High-performance clouds can deliver HPC as a Service
 - High-performance utility computing
- Nearly unlimited application resources
- Instant resource availability
- More efficient use of HPC resources
- Ease of use

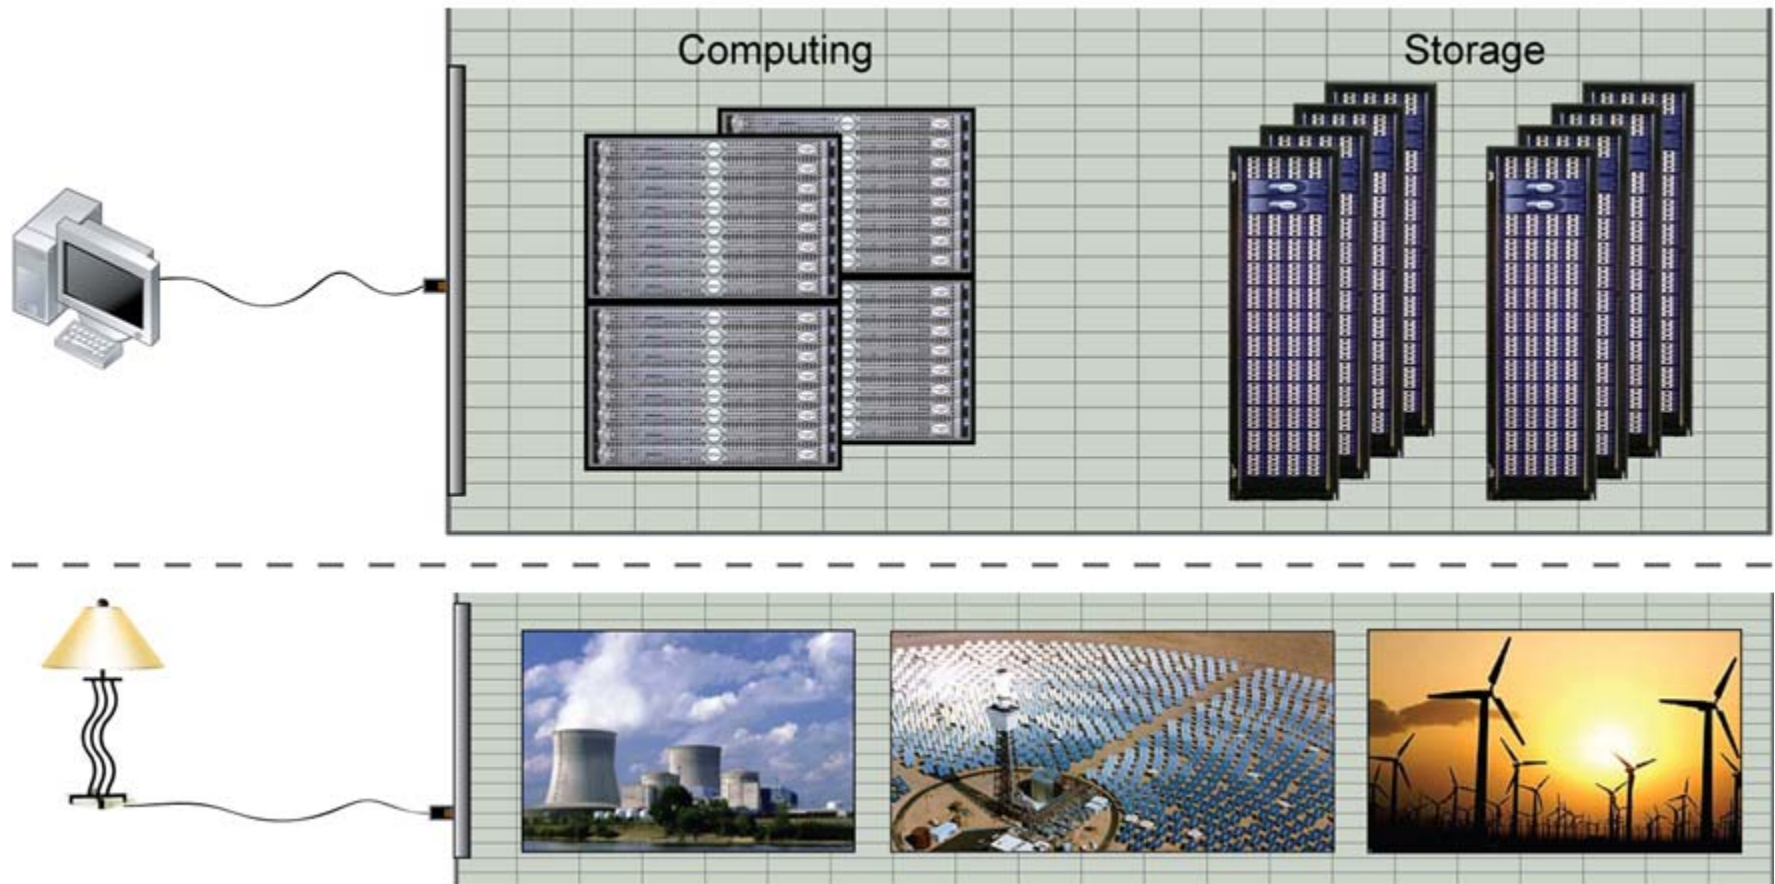


Matching Supply and Demand



The Vision

HPC clouds is not a new concept, and many HPC users have been using clouds for a long time. With emerging technologies and solutions, clouds can serve more users with more applications, and bring to life the concept of utility computing.




Test System Configuration

- **Dell™ PowerEdge™ SC 1435 24-node cluster**
- **Quad-Core AMD Opteron™ 2382 (“Shanghai”) CPUs**
- **Mellanox® InfiniBand ConnectX® 20Gb/s (DDR) HCAs**
- **Mellanox® InfiniBand DDR Switch**
- **Memory: 16GB memory, DDR2 800MHz per node**
- **OS: RHEL5U2, OFED 1.4 InfiniBand SW stack**
- **Platform LSF HPC v7 update5**
- **Platform RTM v2**
- **Platform EnginFrame**




Dell PowerEdge™ Server Advantage

- Dell™ PowerEdge™ servers incorporate AMD Opteron™ and Mellanox ConnectX InfiniBand to provide leading edge performance and reliability
- Building Block Foundations for best price/performance and performance/watt
- Investment protection and energy efficient
- Longer term server investment value
- Faster DDR2-800 memory
- Enhanced AMD PowerNow!
- Independent Dynamic Core Technology
- AMD CoolCore™ and Smart Fetch Technology
- Mellanox InfiniBand end-to-end for highest networking performance



Quad-Core AMD Opteron™ Processor

Uncompromised Design. Unprecedented Performance-per-Watt.



Quad-Core AMD Opteron processors are designed for optimum multi-threaded application performance.

Native Quad-Core Designed for scalability

- Enhanced CPU IPC
- 4x 512K L2 cache
- 6MB L3 cache

Direct Connect Architecture

- HyperTransport™ Technology
- Up to 24 GB/s peak per processor


Floating Point

- 128-bit FPU per core
- 4 FLOPS/clock peak per core

Integrated Memory Controller

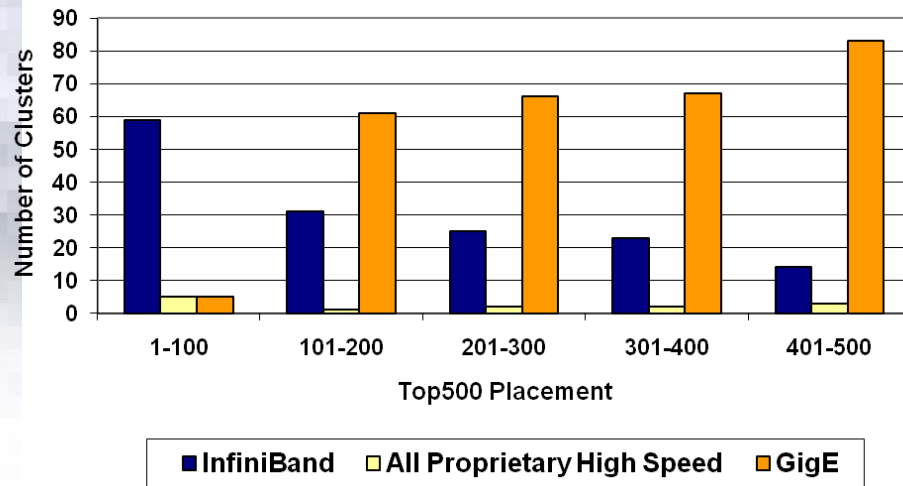
- Up to 12.8 GB/s
- DDR2-800 MHz or DDR2-667 MHz

48-bit Physical Addressing

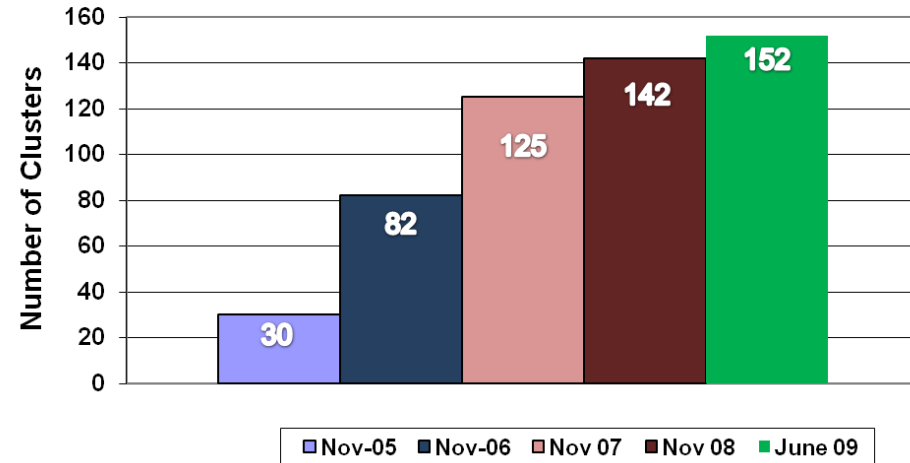


InfiniBand in the Top500 - June 2009

Top500 Interconnect Placement

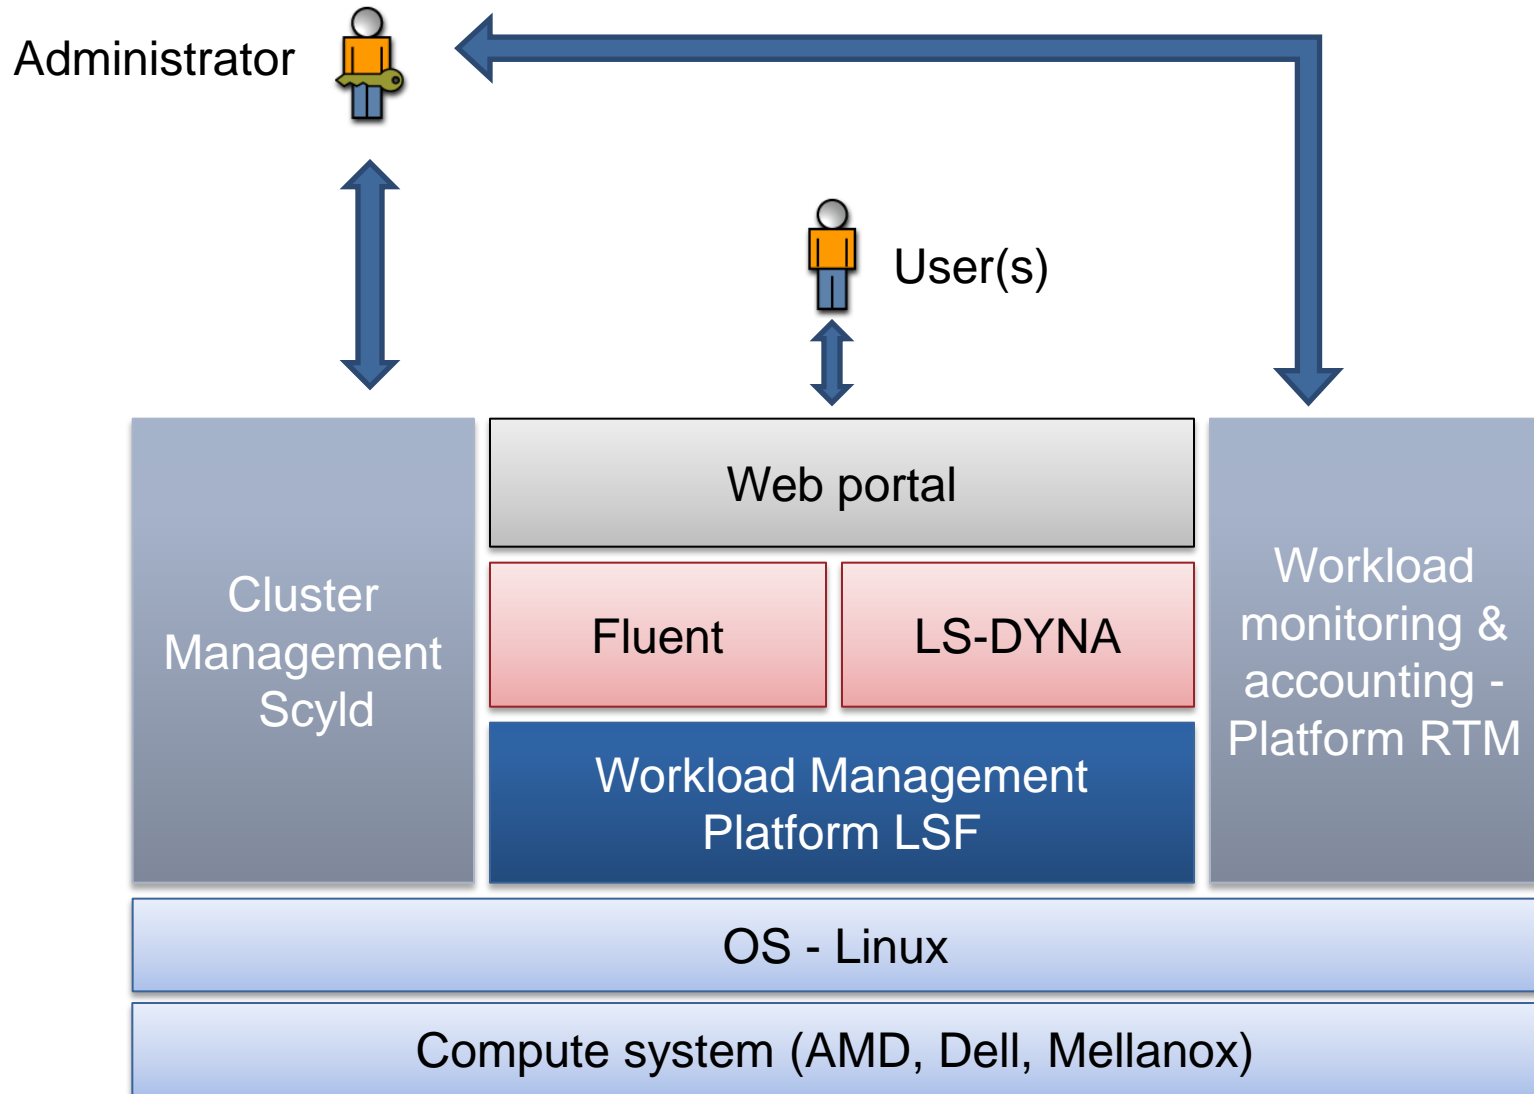


Top500 InfiniBand Trends



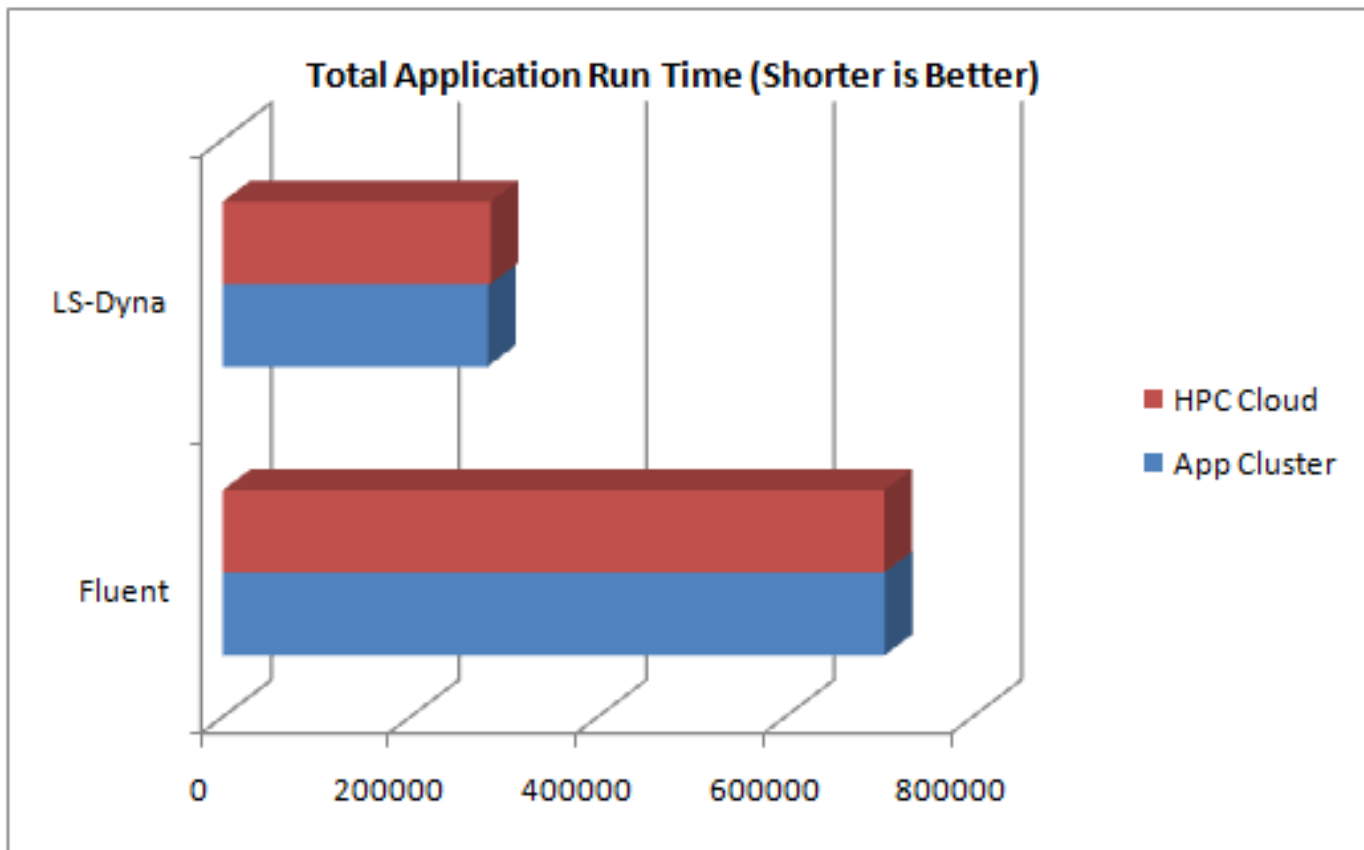
- **InfiniBand is the high performance interconnect of choice**
 - Connecting the most powerful clusters, providing the highest system utilization
 - 40Gb/s-120Gb/s today, 80Gb/s-240Gb/s in 2010
- **InfiniBand is the only growing high speed clustering interconnect**
 - 25% increase since June 08
- **Mellanox 40Gb/s InfiniBand end-to-end the only proven technology on the list**
 - Enable the highest system utilization in the Top10 systems
- **All InfiniBand systems use Mellanox solutions and technology**

HPC in a Cloud Proof of Concept




- **Three provisioned clusters on Osiris**
 - 48-core Fluent system
 - 48-core LS-DYNA system
 - 96-core HPC Cloud system
- **Fluent system:**
 - 16 Fluent (truck_14m) jobs with different size and run time
- **LS-DYNA system:**
 - 16 LS-DYNA (neon.refined.rev01.k) jobs with different size and run time
- **HPC Cloud system:**
 - Rerun those 32 jobs in the HPC Cloud cluster
- **All jobs are submitted by script**
- **The scheduler keeps the cluster as busy as possible**

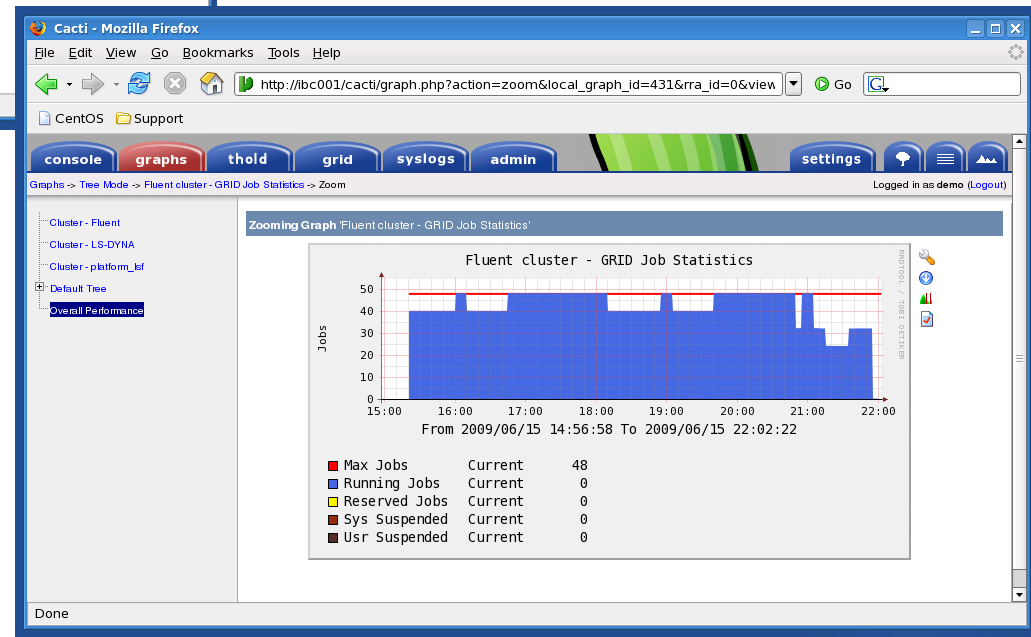
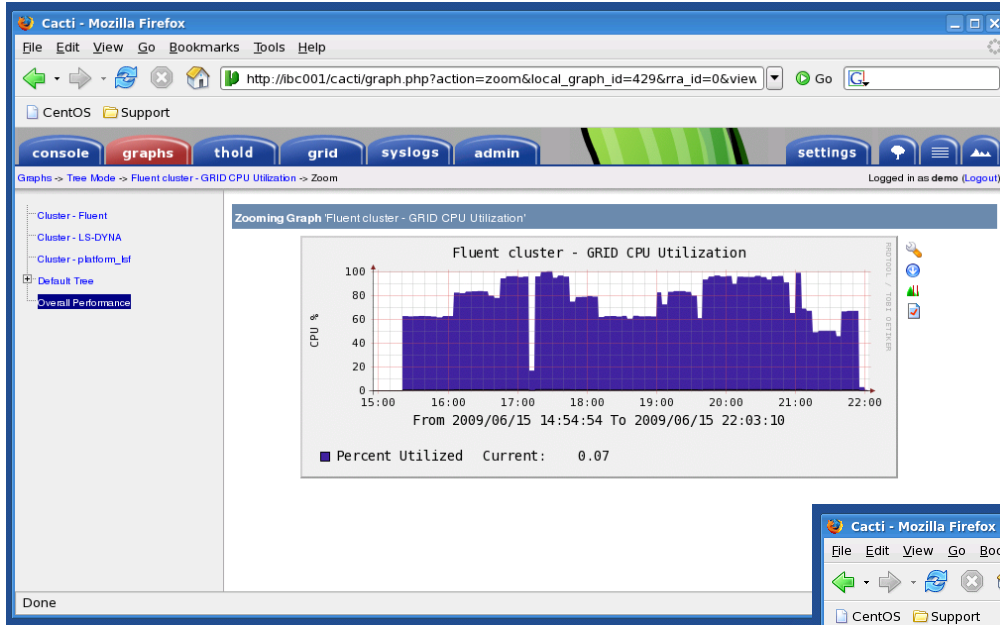
- **The cloud can provide same application run time versus the native approach – enabling HPC in a cloud**



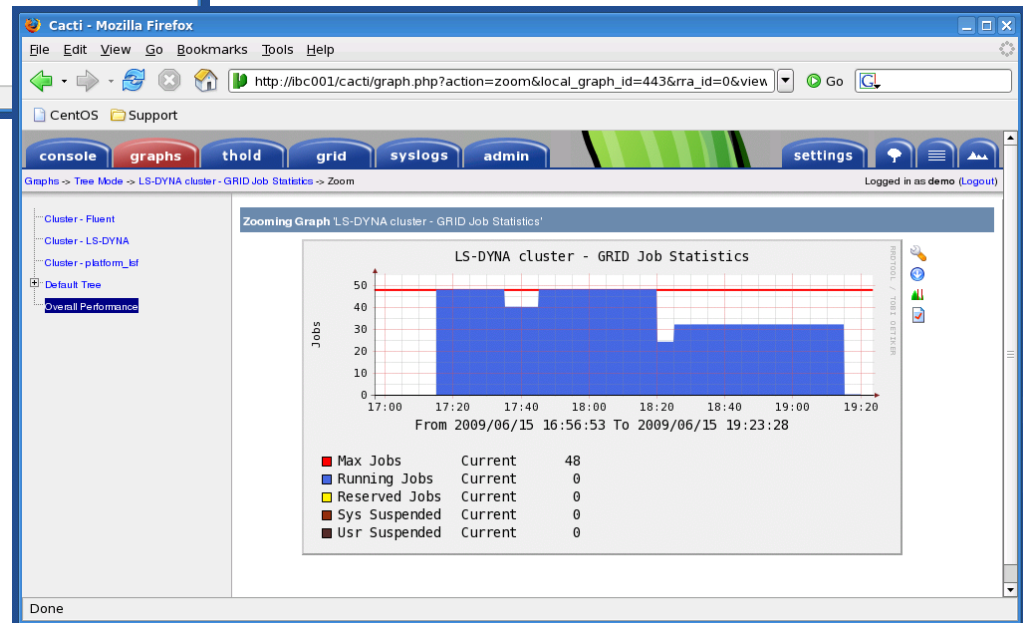

- **System utilization can increase in a cloud**
 - Less idle time due to scheduling or lack of available resources



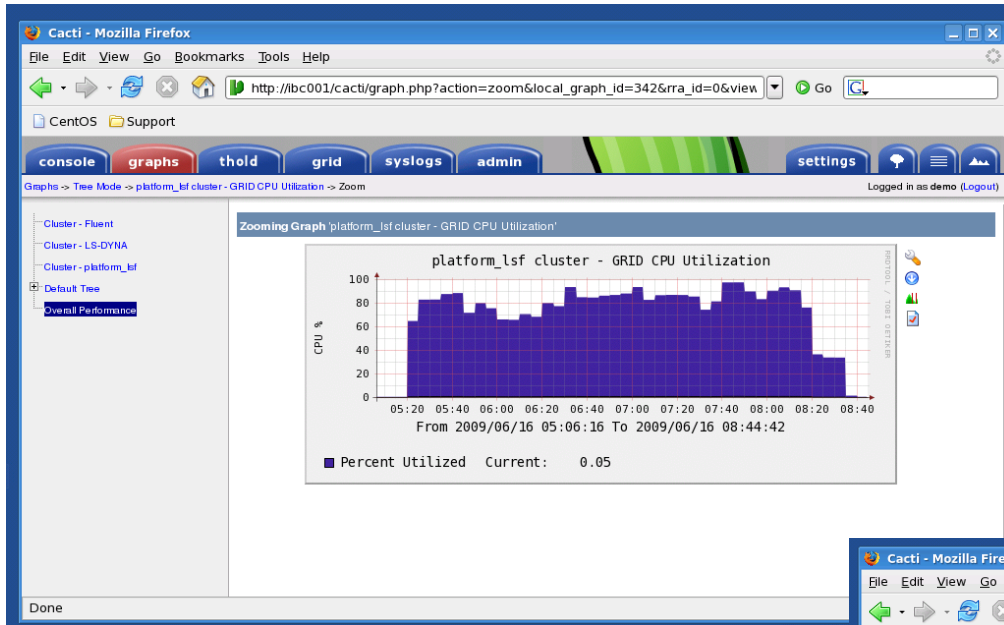
Fluent - CPU and System Utilization Charts



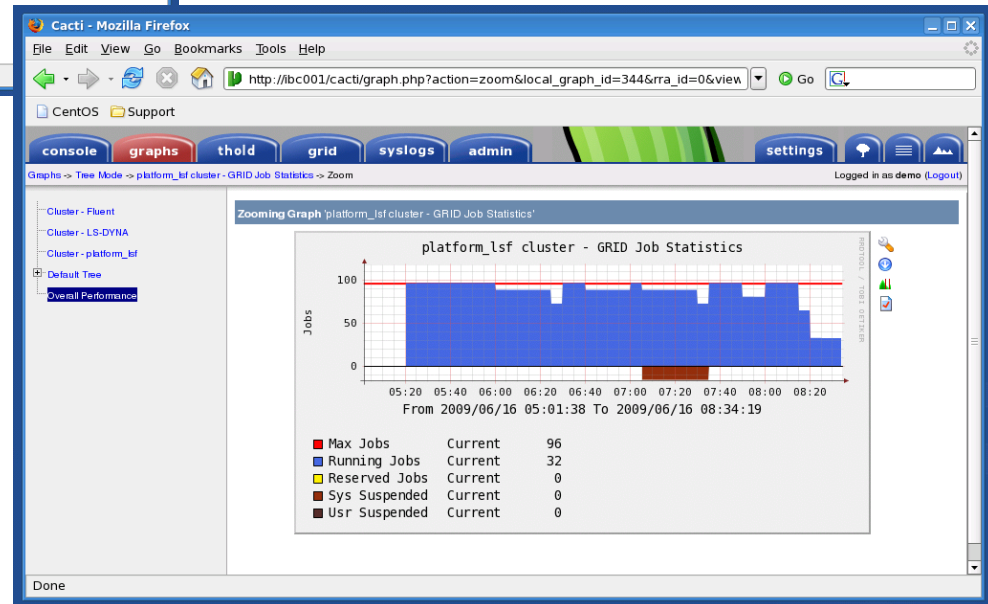
LS-DYNA - CPU and System Utilization Charts



Highest CPU Utilization



Highest System Utilization



- **Cloud Computing provides the ability to deliver Software and Hardware as a Service in HPC markets**
- **System utilization can be optimized in a cloud environment versus a dedicated application system**
- **Organizations should look into providing SaaS and HaaS as the starting point to implement private clouds**

- **Dedicated application environments**
 - Every time application changes, the application runtime environment also changes
- **Running HPC application in VM**
 - Isolated app environment
 - Enable checkpoint/restart
- **Dynamically add more resources into the cloud**
- **Topology awareness for optimized cloud scheduling**

Thank You

HPC Advisory Council

<http://www.hpcadvisorycouncil.com/>



All trademarks are property of their respective owners. All information is provided "As-Is" without any kind of warranty. The HPC Advisory Council makes no representation to the accuracy and completeness of the information contained herein. HPC Advisory Council Mellanox undertakes no duty and assumes no obligation to update or correct any information presented herein