

Discussion of Infrastructure Clouds

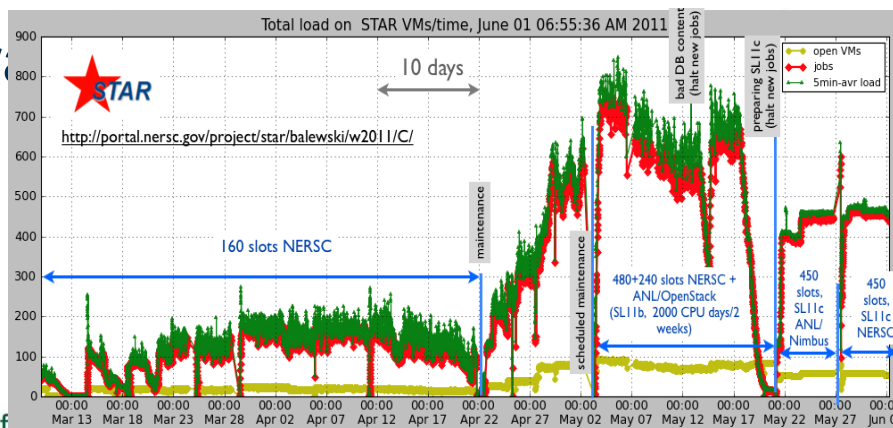
A NERSC Magellan Perspective

Lavanya Ramakrishnan

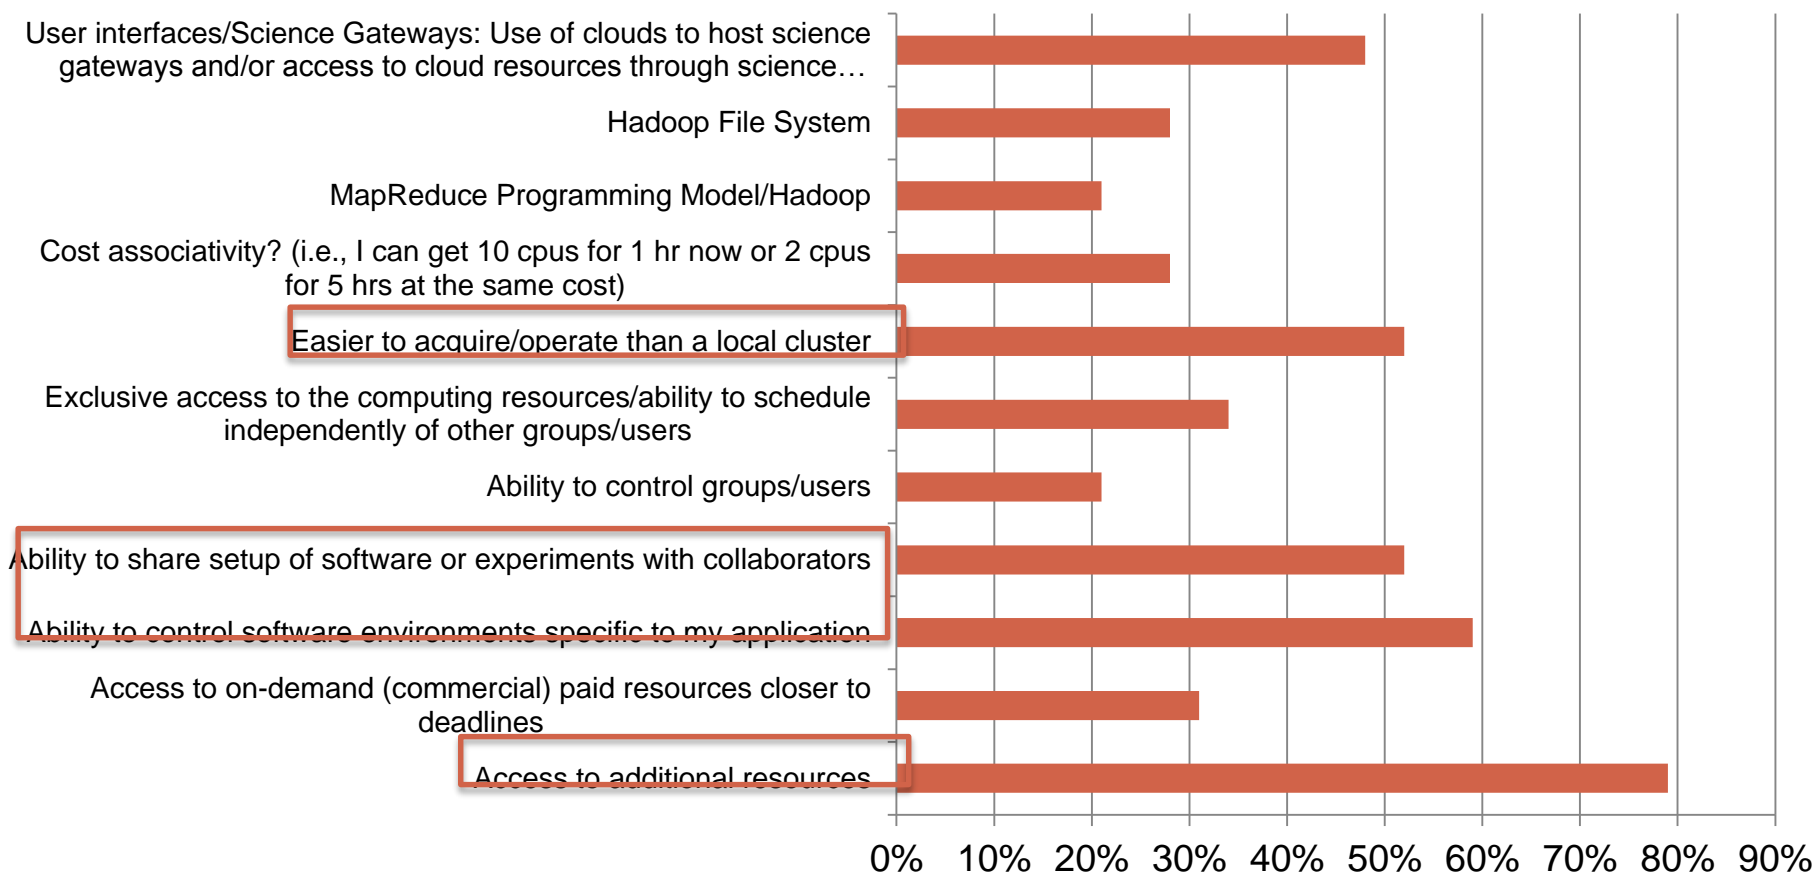
Lawrence Berkeley National Lab

Magellan Applications

- **Magellan has a broad set of users**
 - various domains and projects and workflow styles
- **Example: STAR performed Real-time analysis of data coming from Brookhaven Nat. Lab**
 - first time data was analyzed in real-time to a high degree
 - leveraged the **NERSC system**



Magellan User Survey



Program Office	Percentage
Advanced Scientific Computing Research	17%
Biological and Environmental Research	9%
Basic Energy Sciences -Chemical Sciences	10%
Fusion Energy Sciences	10%

Program Office	Percentage
High Energy Physics	20%
Nuclear Physics	13%
Advanced Networking Initiative (ANI) Project	3%
Other	14%



Challenges of Using Infrastructure Clouds

Order of Significance varies by application

- **Performance, Scalability, Reliability**
 - overheads of virtualizations
 - limit on concurrent VMs
- **Security, Allocation & Accounting**
 - user provided images and root privileges
 - hard to ensure fairness
- **Application Design and Development**
 - image creation and management
 - workflow and data management
 - performance, fault-tolerance need to be considered.

Resource Models

	Traditional Enterprise IT	HPC Centers
Typical Load Average	30% *	90%
Computational Needs	Bounded computing requirements – Sufficient to meet customer demand or transaction rates.	Virtually unbounded requirements – Scientist always have larger, more complicated problems to simulate or analyze.
Scaling Approach	Scale-in. Emphasis on consolidating in a node using virtualization	Scale-Out Applications run in parallel across multiple nodes.

	Cloud	HPC Centers
NIST Definition	Resource Pooling, Broad network access, measured service, rapid elasticity, on-demand self service	Resource Pooling, Broad network access, measured service. Limited: rapid elasticity, on-demand self service
Workloads	High throughput modest data workloads	High Synchronous large concurrencies parallel codes with significant I/O and communication
Software Stack	Flexible user managed custom software stacks	Access to parallel filesystems and low-latency high bandwidth interconnect. Preinstalled, pre-tuned application software stacks for performance

Cloud computing is a business model: can be applied to HPC systems as well as traditional clouds

Questions?

<http://magellan.nerisc.gov>

Lavanya Ramakrishnan
LRamakrishnan@lbl.gov