
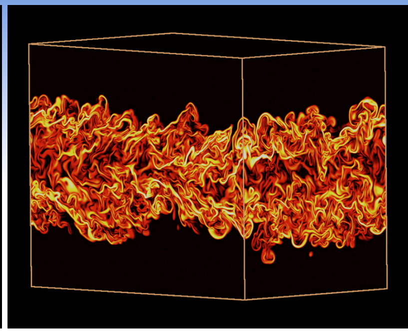
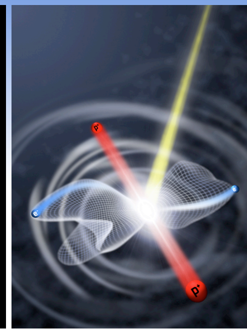
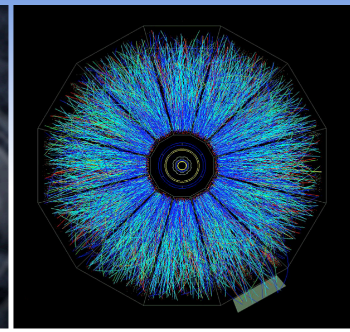
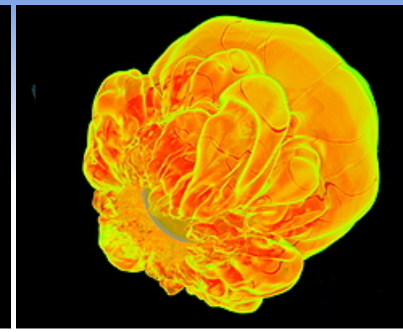

Magellan Project

Jeff Broughton

NERSC Systems Department Head

October 7, 2009

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC

BERKELEY LAB

Magellan Background

- National Energy Research Scientific Computing Center (NERSC)
- Argonne Leadership Computing Facility (ALCF)
- $32M total funding, equally divided between the two facilities
- Funded by DOE under the American Recovery and Reinvestment Act (ARRA)

Magellan Mission

- Deploy a test bed cloud to serve the needs of mid-range scientific computing.
- Evaluate the effectiveness of this system for a wide spectrum of DOE/SC applications in comparison with other platform models.
- Determine the appropriate role for commercial and/or private cloud computing for DOE/SC midrange workloads

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Why is ASCR funding Magellan?

- **ASCR recently hosted a workshop (www.sc.doe.gov/ascr/ProgramDocuments/ProgDocs.html) to assess the role of mid-range computing in the Office of Science and revealed that this computation continues to play an increasingly important role in enabling the Office of Science.**
- **Although it is not part of ASCR's mission, midrange computing, and the associated data management play a vital and growing role in advancing science in disciplines where capacity is as important as capability.**

  – ***Dan Hitchcock, ARRA Project Briefing to ASCAC, Aug 11, 2009***

What is a cloud?

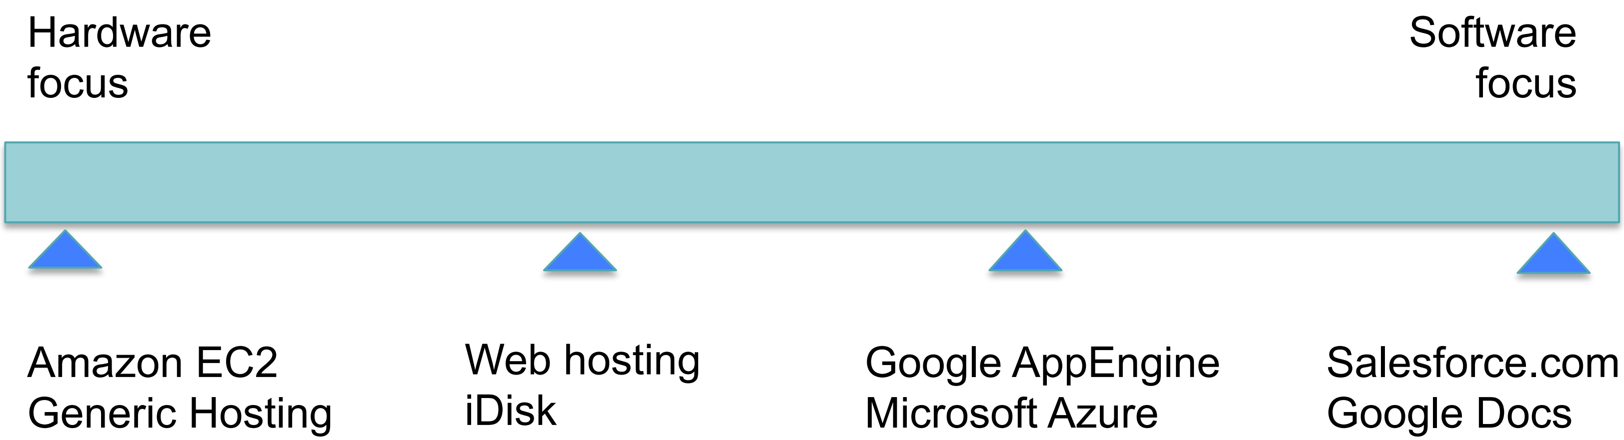

- **Opaque infrastructure**
- **Capacity >> Demand**
- **Available for rent**
- **Self-service**

Perceived User Benefits

- Easy to acquire
- Self-service provisioning
- Pay-as-you-go
- Inexpensive; low up-front costs
- Capacity available “on demand”
- Ability to support “surge” requirements
- Immediate access – jobs run “right away”
- OS/software version selected as required
- Higher availability

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Perceived Management Benefits

- Management by professional administrators (versus research staff)
- Avoids need for local clusters & associated (hidden) costs
- Utilize spare capacity available at large Internet data centers
- Elastic resource model – aggregate uneven demand
- Fractional use of systems – increased efficiency of computing plant
- Transparent swap out of HW for ease of upgrades.
- Location independence – users don't need to care where the systems are.
- Energy efficiency – many clouds located at site of (renewable/green) energy production; follow-the-sun model

What is midrange computing?

- Parallelism < ~1000 cores
- Performance < ~10 Tflops
- System cost < ~$1M

Midrange Computing Alternatives

	Cluster computing	Public cloud computing	Throughput supercomputing
Archetype	Cluster per PI	Amazon EC2	NERSC
Elasticity	None. Fixed peak capacity.	High (excess capacity)	High (load scheduling)
Aggregation method	Ad hoc sharing, or batch system.	Virtualization	Batch system
Acquisition method	Capital purchase	On demand with P.O. or credit card	Allocation
OS	User selected, fixed.	User selectable per image	Single OS across system
Application SW	User selected, licensed and installed.	Typically LAMP, Gnu preinstalled; plus user provided	Preinstalled HPC tools, libraries and apps. Shared licensing
CPU and memory	Can be exact match to application	Tiered service levels.	High bin CPUs. ~2-4GB/core.
I/O and Data Storage	SATA disks or arrays	Highly reliable block or blob storage; metered performance	Multi GB/s global shared file system; Archival storage
Interconnect	1Gb Ethernet; sometimes 10GbE or IB	1Gb Ethernet; moving to 10GbE	Low latency; high bandwidth, e.g. Cray XT or IB
Support model	Self-support	Basic support for configuration and supplied software. User application not supported.	Fully supported; consulting services
Security	Depends on admin	Sophisticated between images (e.g. image isolation; VPNs)	Sophisticated (e.g. intrusion detection; exploit avoidance)
Best application fit	Computation tightly bound to local data acquisition	Loosely coupled applications with small datasets.	Loose or tightly coupled applications; high I/O rates; high concurrency
Cost	Highest of three. OpEx dominates CapEx, but no economy of scale.	$0.10/core-hour; $0.10-0.17/GB network transfer; $0.10 per GB/mo + $0.10/Mops data storage	<$0.14/core hour. Includes storage, network use, and support (est. <$0.04 in 2010)
Issues	Hidden costs; lost of focus on research	Multiple large images may not be "on demand"	Part metric discourages use for small-to-medium jobs.

Areas of Investigation

- Suitability of alternative delivery models (HaaS, DaaS, SaaS)
- Ease-of-use
- Computational efficiency of significant DOE/ SC applications across alternative platforms
- Cost efficiency analysis including energy efficiency, TCO and utilization models.
- Applicability of cloud computing APIs to science apps
- Practicality of multi-site clouds

Hardware as a Service (HaaS)

- Provide a cloud computing service for data-intensive, moderately parallel jobs. Explore:
  - Batch, interactive & virtual private clusters
  - On-demand division of workload; overflow provisioning
  - Virtualization & fixed OS
  - Support for science apps, build & test, infrastructure (web, twiki, gateways, etc.)

Data as a Service (DaaS)

- Provide a shared storage system accessible to local and remote clouds using WAN access to GPFS and GridFTP over ESnet.
- Is it practical to provide a remote data storage service at high bandwidth?
- Does this permit applications to move across sites transparently?

Software as a Service (SaaS)

- Can some popular mid-range applications be provided as a service within the cloud, and thus make scientists more productive?

Ease-of-use

- What should a science cloud really look like?
- Would changes to commercial clouds make them more appropriate for mid-range scientific computing?
- What are the requirements for a private cloud to support mid-range scientific computing?
- What factors make a cloud easy to use?
- How does a science cloud differ from capacity/throughput supercomputing?

Workload Efficiency

- What part of the DoE/SC workload can be supported within an existing commercial cloud model and which is most appropriate?
- What is the efficiency of computing in the cloud and how does that depend on workload characteristics?

- Identify a small number of applications as cloud incubation projects
  – Signifcant DOE/SC workloads
  – Presenting a range of challenges for clouds and clusters
- Instrument cluster to accurately characterize the concurrency, communications patterns and I/O of applications to select candidates to port to clouds.
- Measure the performance of real science workloads operating in multiple environments: batch with IB and/or Ethernet, virtualized systems, commercial cloud.

Cost Efficiency

- How cost effective are commercial clouds relative to privately owned clouds for various workloads? Versus DoE's current ad-hoc approach to mid-range computing?
- What are the tradeoffs between utilization and response time? How does that effect the economics and usability of commercial and private clouds?
- What are the appropriate cost metrics for a science cloud?

Use of Cloud Computing APIs

- Integrate cloud API's for accessing data with DoE data storage facilities
- Are cloud storage APIs applicable to typical scientific applications? Is this an effective way to share data?
- Is MapReduce/Hadoop applicable for data parallel applications (such as HEP and Genomics)?

Multi-site Clouds

- Can a single DoE front-end to a cloud provide a convenient mechanism for accessing both DoE and commercial clouds?
- Can ANL and NERSC clouds provide high availability by allow job queues at one site to be redirected to another?
- Can DoE clouds run at different security levels within the same logical cloud?
- Can ALCF and LBL efficiently join Hadoop clusters?

ANI Requested Activities

- The Advanced Network Initiative is an upgrade of Esnet to 100Gbit/s, funded by ARRA
- Provide a demonstration vehicle for the 100Gbit/s link
- Support OSG and ESG ANI/ARRA projects

In conclusion…

- NERSC is at the forefront of scientific cloud computing
- Please participate in the cloud test bed activities
- Provide feedback on how to build a cloud that best suits the needs of scientists

 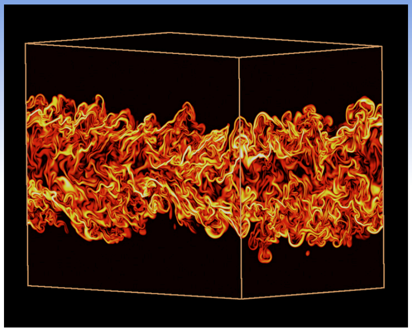 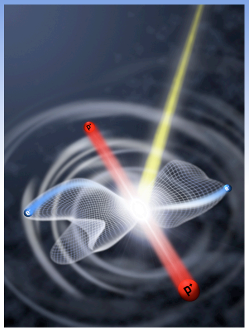 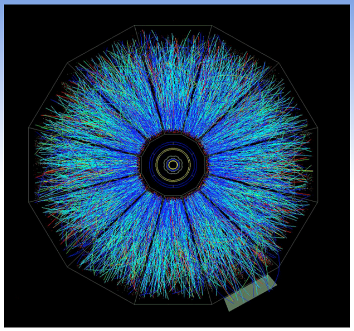 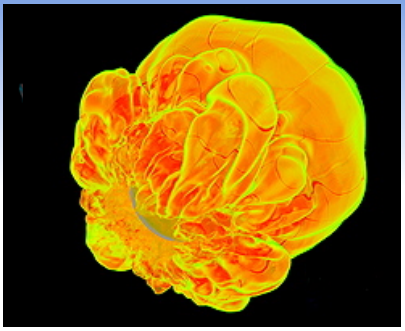

Thank you!

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

NUG Town Hall Meeting Questions

- Are you interested in using cloud computing?
- What benefits do you expect to receive?
- Have you actually tried a commercial cloud?
- What comments do you have on your experience?
- Do you have a specific project that you would like to try using cloud computing?