OpenGridForum
OPEN FORUM | OPEN STANDARDS

Open Grid Forum's Cloud Perspective

July 7, 2009

Craig A. Lee
President, OGF, www.ogf.org

Vast Array of Application Scenarios Are Driving Evolution

OpenGridForum

- Integrated Physical and Computational Systems
- Dynamic Data-Driven Application Systems (DDDAS)

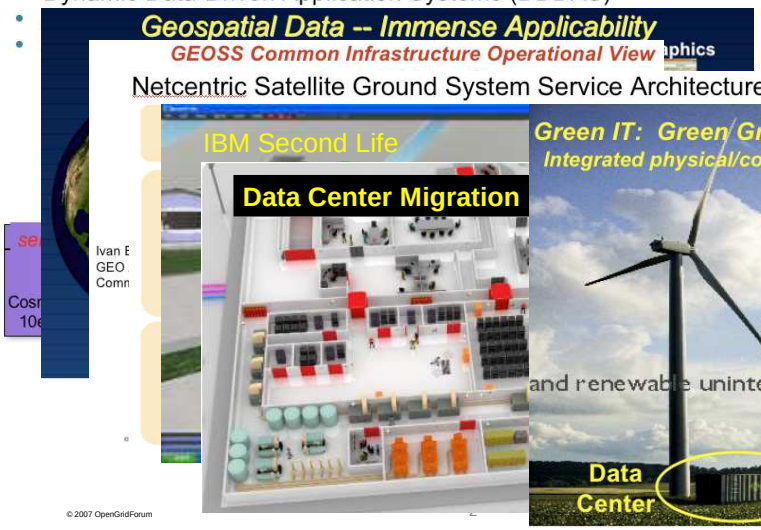

So How Did This Lead up to Clouds?

OpenGridForum

- Distributed Computing has been around since the first packet flowed in 1969
  - Since then many "flavors" of distributed computing tools and infrastructures have been built: DCE, CORBA, PVM, Clusters, Grids, SOAs, ...
- Capabilities were driven/enabled by the installed infrastructure
- World Wide Web led to explosive growth in network use
  - Web hosting/search, e-commerce, social computing have driven enormous growth of data centers
  - Google, Amazon, e-Bay, Facebook, YouTube, iTunes, SecondLife, …
- Data Centers provide enormous compute and data storage capacity
  - Every third processor built goes into a data center
- Data Centers have undeniable economies of scale
  - Google can buy processors for 1/7th of what you can
- Cloud computing enables data center operators to sell capacity
  - The cloud model enables clear *business models*
- Other names:
  - utility computing, internet computing, …

www.ogf.org

3

Cloud Computing -- What is it?

- A broad term used to denote *abstraction* and *virtualization* at any of several different system layers
  - "Outsourcing" of hardware, system environment, or services
  - Things just run "in the cloud", i.e., somebody else's data center
- Generally from a single provider through a very simple API
  - Simple API eases adoption at the cost of insight and control
  - Effective business model for provider to "sell" virtualized, back-end data center resources

Application Level	• *Software as a Service (SaaS)* • Build an application from pre-defined services • Example: Salesforce.com
Platform Level	• *Platform as a Service (PaaS)* • Acquire a set of hosting environments • Example: Google App Engine (Python)
Infrastructure Level	• *Infrastructure as a Service (IaaS)* • Acquire a set of machines you can login to • Example: Amazon EC2

Key Cloud Characteristics

- On-demand self-service
- Ubiquitous network access
- Location independent resource pooling
- Rapid elasticity
- Pay per use

Peter Mell and Tim Grance - NIST, Information Technology Laboratory

www.ogf.org

5

Expected Cloud Benefits?

- Commodification of compute infrastructure
  - General infrastructure that can support many programs, functions
- Improving server utilization
  - Flexibility in mapping work to servers
- Managing surge requirement with a pool of common resources
  - Sizing system for the average case, rather than the worst case
- Improving reliability
  - Easier fail-over between servers
- Greener IT
  - Reduce energy costs through consolidation, improved utilization & moving work to the cheaper/greener energy
- Reduced Costs
  - “Right-sized” commodity infrastructure with improved utilization

www.ogf.org

6

How Do Grids and Clouds Relate?

- Opinions vary … but …
- Grids came out of "big science" and the desire to collaborate in a *federated environment*
  - Managed sharing of resources
- Clouds came out of industry and the desire to dynamically provision resources *in the cloud*
  - Simple APIs for using abstracted or virtualized resources
  - Economies of scale in the data center
  - Aka, utility computing, internet computing, …
- "Grids are an access model; Clouds are a business model"
  - Chris Smith, Platform Computing, OGF VP Standards
- Distributed applications need and can use capabilities being developed under both rubrics of *grid* and *cloud*
  - There is no real *grid vs. cloud dichotomy*

A Possible Way to Access Cloud Services

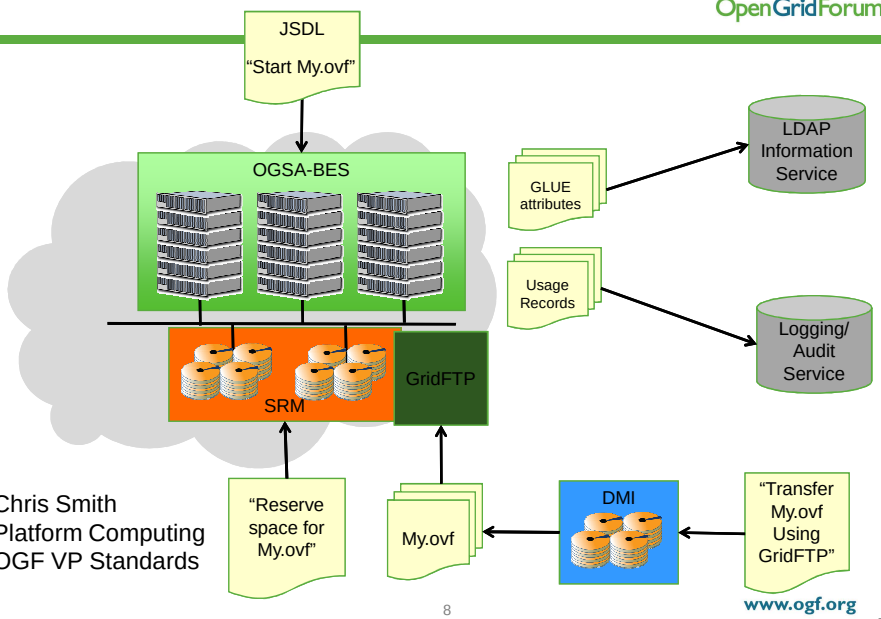

Chris Smith
Platform Computing
OGF VP Standards

Referenced Standards

- Web Services Agreement Specification (WS-Agreement)
  - http://www.ogf.org/documents/GFD.107.pdf
- GLUE Schema
  - http://www.ogf.org/documents/GFD.147.pdf
- Lightweight Directory Access Protocol (LDAP)
  - http://tools.ietf.org/html/rfc4510
- Storage Resource Manager Interface (SRM)
  - http://www.ogf.org/documents/GFD.129.pdf
- Data Movement Interface (DMI)
  - http://www.ogf.org/documents/GFD.134.pdf
- GridFTP
  - http://www.ogf.org/documents/GFD.20.pdf
- Open Virtualization Format Specification (OVF)
  - http://www.dmtf.org/standards/published_documents/DSP0243_1.0.0.pdf
- Job Submission Description Language (JSDL)
  - http://www.ogf.org/documents/GFD.136.pdf
- Basic Execution Service (BES)
  - http://www.ogf.org/documents/GFD.108.pdf
- Usage Record (UR)
  - http://www.ogf.org/documents/GFD.98.pdf

No Shortage of Challenges

- Data access and interoperability
  - Must be done at the application domain level, by the domain users
- Security
  - Different models will expose different security threats
- Reliability
  - Managing redundancy, live migration, etc., across the infrastructure
- Frameworks
  - How to manage sets of resources, e.g., VMs and VOs?
- Performance management
  - What job mix needs to be supported, e.g., e-commerce, HPC, transactional, database, data streaming?
- Costing models
  - How to compare your own infrastructure costs with a cloud providers?

An Evolving Infrastructure

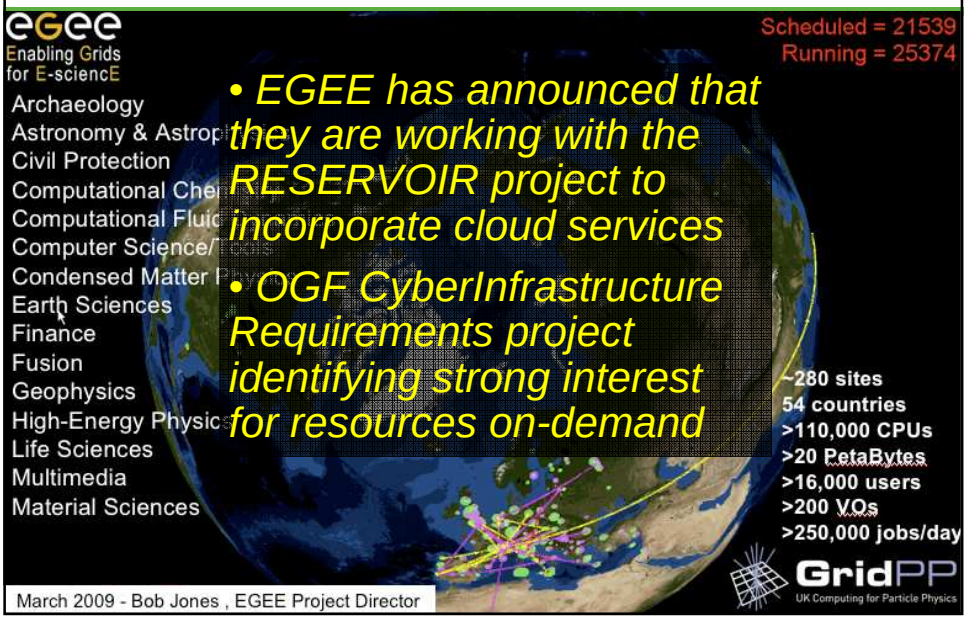

Deployment Models Blur the Boundaries of Grids and Clouds

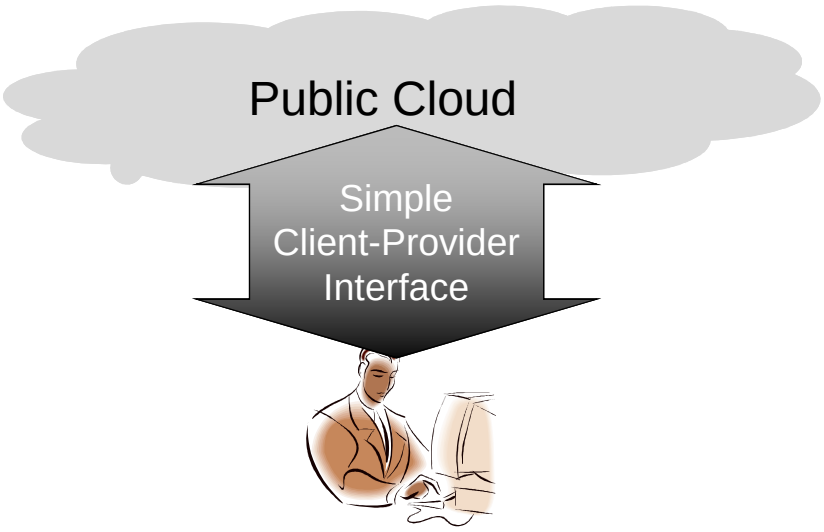

Deployment Models Blur the Boundaries of Grids and Clouds

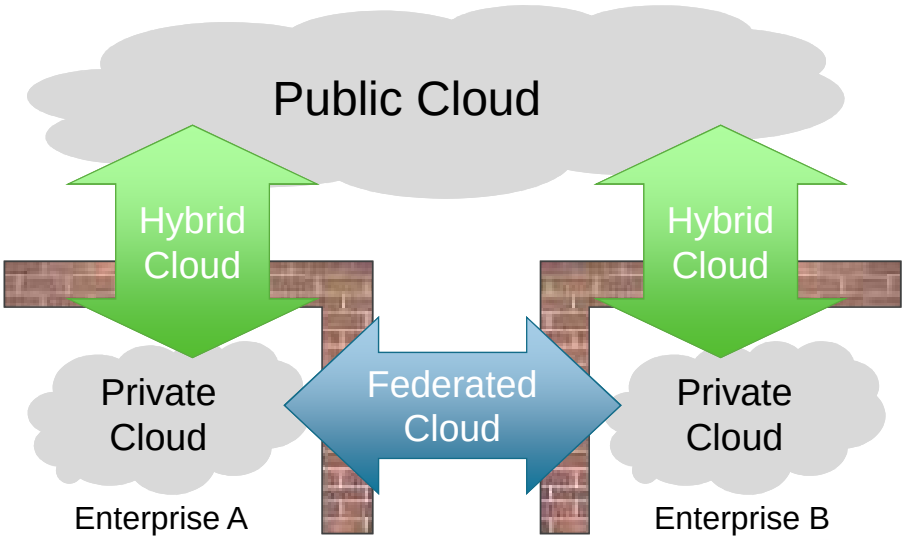

Public vs. Private Cloud Issues

- Cost & Cost Predictability
- Users expect to monitor & manage "their" infrastructure
  - Will a public cloud provider expose enough information for a client to troubleshoot when something goes wrong?
- Security & Privacy
  - You can store encrypted data in a clouds, but can you compute on it?
- Regulation
  - Physical location of data
  - Long-term audit trails (15-20 years)
- Individual vs. Corporate Requirements
  - Corporate use of public clouds may entail legal & contracting overheads
  - Ease of use and quick provisioning may tempt individuals to ignore corporate procedures
  - Trade-off between quick results and risk exposure
- Internal IT departments may want to offer their own "seed cloud"
- Interoperability & portability between private and public clouds

The OGF Open Cloud Computing Interface

- OCCI focused on IaaS Cloud API
  - Requirements and Use Case documents also underway
- Goal was the creation of a simple and RESTful API
  - Simple means ~15 commands
  - ***Slim -- very extensible!***
- OCCI collaborating with DMTF, SNIA and SCRM

http://www.occi-wg.org

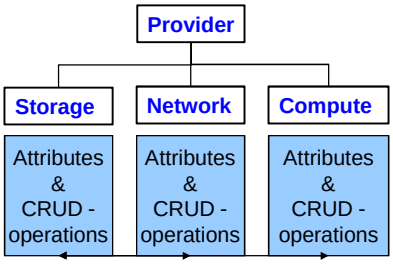

- WG chairs represent whole cloud community
  - Academia – Andy Edmonds (SLA@SOI - Intel)
  - Industry – Thijs Metsch (Sun Microsystems)
  - End-Users – Sam Johnston (Australian Online Solutions)
  - Service Providers – Alexis Richardson (RabbitMQ & CohesiveFT)
- Interest has been huge
  - 160 members on mailing list
  - Four providers and two projects to implement standard

Links

www.ogf.org

Green IT/Grids/Clouds

- Energy becoming defining issue for Data Centers
  - Beyond cost: location feasibility
- Energy Management Architectures are Sensor & Control Systems
  - Environmental Monitoring
  - Decision Making
  - Enacting Responses
- Control Mechanisms include:
  - Scheduling
  - Voltage-Frequency Scaling
  - Powering racks up-down
  - VM placement & migration
- Clouds can be used to enforce energy policy
  - By abstracting the infrastructure, clouds can transparently manage workload on the back-end
    - Consolidate jobs on servers
    - Move work to where the green power is
- Several projects underway
  - Reservoir (EU), OpenNebula (EU), GreenNet (INRIA), GreenLight (NSF), Low Carbon ICT (UK)

More data centers are just not possible in NYC or London

An Energy Management Architecture

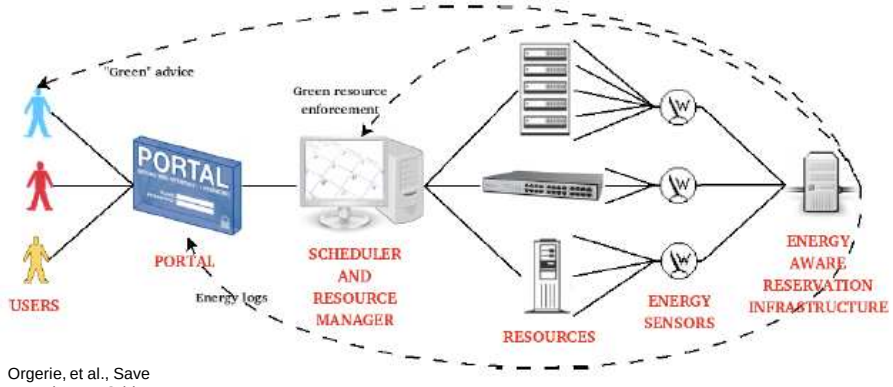

Orgerie, et al., Save Watts in your Grid: Green Strategies for Energy-Aware Framework in Large Scale Distributed Systems. 14th IEEE ICPDS, Dec. 2008.

- CO2 Working Group Starting
  - Reporting interface between physical infrastructure management and workload management

Pushing for Cloud Coordination

- SATCCI Meeting, March 23, DC
  - federalcloudcomputing.wik.is/index.php?title=March_23%2C_2009
- SCRM telecons
  - OGF, DMTF, SNIA, OMG, OCC, NIST, IBM, Cisco, Sun, Fujitsu, attending calls
  - Organizations posting cloud activity summaries on SRCM wiki
  - On July 13, 2009, forge.gridforum.org/sf/wiki/do/viewPage/projects.scrm-wg/wiki/CloudStandardsCoordination will be promoted to *cloud-standards.org*
- Cloud Standards Summit, July 13, DC
  - federalcloudcomputing.wik.is/July_15%2c_2009
  - Originally scheduled for July 15 but moved to accommodate the following:
- Cloud Computing Symposium, July 15, DC
  - Hosted by National Defense University, www.ndu.edu/irmc/ilss/ilss_events.html
  - Vivek Kundra (US Federal CIO) to keynote
  - Cloud Standards Summit has an outbrief slot
- CloudWorld
  - Co-located w/ NGDC and OpenSource World, August 12-13, San Francisco
  - OGF personnel on CloudWorld Program and Advisory Committees
  - Chris Smith (OGF VP Standards) to present on behalf of OGF Enterprise

Economic Realities

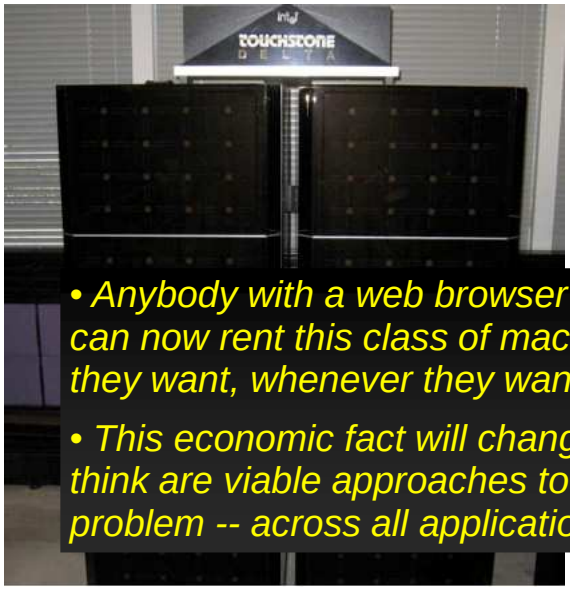

- Intel Touchstone Delta Machine Installed at Caltech, May 1991
  - 512 nodes
  - >4 GB of memory
  - #8 on first Top 500 List, June 1993
  - Decommissioned 1998

- *Anybody with a web browser and a credit card can now rent this class of machine for as long as they want, whenever they want*
- *This economic fact will change what people think are viable approaches to solving their problem -- across all application domains*

www.ogf.org

Upcoming OGF Events!

- OGF-27
  - Banff, Alberta, Canada, Oct 12-16, 2009
  - With IEEE Grid 2009 & Cybera/CANARIE Natl Summit
  - Combined “Grid to Cloud” & “HPC in Cloud” Workshops

- OGF-28
  - Munich, Germany
  - March 8-12, 2010

Take-Home Message

- Wide spectrum of applications can use the capabilities being developed under the rubrics of *grid* and *cloud*
- OGF will energetically pursue -- in collaboration with the broader community -- all useful computing technologies of interest to our stakeholders
- OGF's ultimate goal is the integration of the best technologies available to solve the problems at hand
  - ***We must coordinate and collaborate***

Thank You

lee@aero.org