

MAGIC Meeting
February 2, 2011, 2:00-4:00

Attendance:

Linda Akli	SURA
Bob Bohn	NIST
Rich Carlson	DOE
Shantenau Jha	LSU
Dawn Leaf	NIST
Miron Livny	OSG
Mike nelson	Georgetown U.
Larry Peterson	Princeton
Ruth Pordes	OSG
Don Riley	U. of Maryland
Mike Seablom	NASA
Alan Sill	OGF
Bill Turnbull	DOE
Wendy Wiggen	NCO

Action Items

1. Any agency interested in helping NIST document their business case uses of cloud computing, please contact Alan, Bob or Dawn. Alan.Sill@ttu.edu , robert.bohn@nist.gov , DLeaf@doc.gov
2. Divy Agrawal is involved in the workshop being organized at NSF and will send details.

Proceedings

This meeting of MAGIC was coordinated by Rich Carlson of DOE and Wendy Wiggen of the NCO.

Business Use Cases for Cloud Computing: NIST, Alan Sill :

NIST has been charged with developing a roadmap process for the US Federal approach to cloud computing. An overall roadmap will be delivered by end of this fiscal year. NIST has, to date, held two workshops in 2010 (see links below). There was some attempt to engage with non-US entities. Preference is to use international agreed upon standards. Dawn's talk (see link) explains where the charge came from. Default purchasing for Federal agencies is now based in cloud computing. Outreach to industry and standards bodies and branches of the agencies is taking place. Five working groups were established.

The process is driven by existing business use cases. NIST is approaching Magic to ask for help and independent involvement by the Science agencies to describe their business use cases.

There has been a decade of valuable work on grid standards; now the grid is interfacing with the cloud. Nebula, Magellen, and Nimbus projects are instantiating cloud computing. One purpose of this briefing is to draw attention to the NIST wiki and solicit support for the effort

Bob Bohn indicated that Vivek asked NIST to create standards. There are 5 working groups. You can get involved at any level you want. NIST has plans for April 7-8 workshop in Gaithersburg. As

a starting point NIST needs names and agencies willing to let their cloud computing projects be documented. Cloud to Grid and Grid to Cloud transmission space is the focus. Distributed computing has a large set of standards that NIST believes can be reinterpreted for the cloud environment. NIST seeks to document agency projects to help define this new environment.

AI: Any agency interested in helping NIST document their business case uses of cloud computing, please contact Alan, Bob or Dawn. Alan.Sill@ttu.edu , robert.bohn@nist.gov , DLeaf@doc.gov

Cloud Computing Forum & Workshops I and II: (Next workshop Apr. 7-8, 2011)

<http://www.nist.gov/itl/cloud.cfm>

<http://www.nist.gov/itl/cloud/cloudworkshopii.cfm>

Cloud Computing Forum & Workshop II Standards Panel:

http://www.nist.gov/itl/cloud/upload/Standards_Panel.pdf

For the full details of the NIST briefing, please see:

<http://www.hpcc.ttu.edu/asill/ogf/nitrd/MAGIC-Dec2010-OGF-Standards-Intro-Web.pdf>

Cloud Computing Research: Divy Agrawal agrawal@cs.ucsb.edu

Divy Agrawal of UCSB described his research in a talk on “From a Virtualized Computing Nucleus to a Cloud Computing Universe: Data Management in the Cloud”.
Focus: Data Management in the Cloud

- We have moved toward defining concepts of cloud computing
- Infrastructures have moved from enterprise owned to commodity and consequently new players have emerged
- Now the challenge is to take existing data management components and make them cloud friendly; data bases are the bottleneck and need to be made as flexible and scalable as the infrastructure
- Economics of Internet users – over vs. under provisioning- There is a heavy penalty for under-provisioning
- Economics of cloud computing – Pay by use instead of provisioning for peak
- We are good at designing infrastructure that is optimized for given conditions – how do we design software to adapt to this?
- Grid computing is similar to cloud computing - both deal with computing as a commodity
 - Virtualization has helped grid computing evolve into cloud computing
 - Alan Sill – cloud computing is not a reinvention or revolution but an evolution from grid computing
 - Cloud computing builds on what has been done in data management. We need to leverage work in that area
 - Bob Bohn asked how elasticity is handled in the grid? A: Grid was primarily for scientific applications that worked on the “reservation” model. Grid didn’t look at 24x7 service that would be going up and down at will.

- Alan Sill – tools have been invented to handle “on-demand” requests and vary from one organization to another. How do we bring those tools together under cloud computing?
- We are facing a new generation of application developers (i.e. Facebook or Animoto.com) that need to be hosted and yet can grow very quickly
- Can add servers, but relational databases become the new bottleneck
- Do away with RDBM and move to key value stores?
- Makes it difficult for application providers; limits applications consistency and flexibility
- How to build programs to manage data when web-based systems demand a different model from the relational database.
- Key value stores - must do updates one at a time – failures can happen – use system of persistent queues that will eventually catch up
- Quote from James Hamilton of Microsoft: eventual consistency (as opposed to strong consistency) comes at the price of programming model complexity
- System Metadata must be consistent
- Data Fission – making RDBMS cloud friendly: metadata manager and owner transaction manager – want to be able to consolidate and transfer loads around to provide elasticity – yet minimally impacting the users and do it on the fly- main point – how to provide elasticity in the database layer to maintain speed of transactions and minimize cost
- Virtualized Nucleus to Cloud computing Universe is focus of current work

AI: Divy Agrawal is involved in the workshop being organized at NSF and will send details.

OGF data standards for structured and unstructured data access are summarized at the link below:

<http://www.ogf.org/gf/page.php?page=Standards::DataArea>

The INFO-D, DAI, and DF DL-related standards are, I believe, relevant to the NIST use cases for clouds.

The new OGF Open Cloud Computing Interface (OCCI), which is emerging as a general-purpose interface to all cloud layers and abstractions is described at:

<http://occi-wg.org>

Finally, the Cloud Data Management Interface (CDMI) standard from SNIA is relevant (SNIA is an OGF partner):

http://www.snia.org/tech_activities/standards/curr_standards/cdmi/

Note that logging, queues, serialization and the role of data containers are addressed explicitly within CDMI and may be of interest.

Grid Roundtable

- DoE is looking at Magellan and cloud computing and standards issues that need to be addressed
- NASA Nebula cloud is now at 2 of 10 facilities; there is a new deployment at NASA Goddard; database issues have not been thought through, ours have been based on computation; Nebula is having an issue on programmatic side whether cloud services should be purchased outside or provided within. The NASA rover team is using Amazon; the concept of “build it and they will come” has so far not been working;

Agency Meetings of Interest

- AFCEA – Meeting in March, cloud computing with a focus on High End Computing.
<https://www.signup4.net/Public/ap.aspx?EID=HIGH37E&TID=LSwScFEn95FXu4T%2bIXNkow%3d%3d>
- OGF – next event in Taiwan in March; looking at grid to cloud workshop in US perhaps in conjunction with Teragrid meeting, Salt Lake City; smaller events April 16 Boulder Cloud Standards Interoperability plug Fest – remotely available; May, in Boulder, Cloud Standards summit; looking for a workshop in late spring early summer
- Information on Grid workshop standards and grid- CCGrid 2011 webpage
<http://www.cloudstrategypartners.com/InterCloudGrid11.htm>

Next MAGIC Meeting: March 2, 2:00-4:00, NSF, Room TBD