



MAGIC Meeting Minutes

January 9, 2013

Attendees

GJ Ahn	Arizona State U.
Jim Basney	NCSA
Rich Carlson	DOE/SC
Dan Gunter	LBL
Dan Katz	NSF
Miron Livny	U. of Wisconsin
Mark Luker	NCO
David Martin	Northwestern U.
Stuart Martin	ANL/OGF
John McGee	RENCI
Grant Miller	NCO
Reagan Moore	RENCI
Ruth Pordes,	OSG/ANL
David Proctor	NSF
Martin Swaney	U. Delaware
Von Welch	Indiana U.

Action Items

1. Reagan Moore will send MAGIC use case examples from the Federal community
2. Grant Miller should include the URL for MAGIC Wiki access on the MAGIC agenda.

Proceedings

This MAGIC Meeting was chaired by Dan Katz of the NSF and Rich Carlson of DOE/SC. Ruth Pordes led a presentation and discussion on Distributed Storage for Big Data Systems.

Distributed Storage Management: Initial Discussion: Ruth Pordes

Ruth Pordes introduced the topic of distributed storage management which, we anticipate, will be discussed on a continuing basis in MAGIC.

Our Goal: Understand needs for robust, fault tolerant storage management systems, particularly in a distributed context - as opposed to part of a single exascale system. Identify any gaps in available systems and technologies and areas where research will be of benefit.

Discussion identified that:

FOR OFFICIAL GOVERNMENT USE ONLY

c/o National Coordination Office for Networking and Information Technology Research and Development

Suite II-405 · 4201 Wilson Boulevard · Arlington, Virginia 22230

Phone: (703) 292-4873 · Fax: (703) 292-9097 · Email: nco@nitrd.gov · Web site: www.nitrd.gov

- How do we know what the science communities need?
- We need a common vocabulary for discussion
- To identify gaps you need to know current status

How do we know what the science communities need?

Is storage different from data distribution and access?

Discussion identified that:

- There is a boundary between storage and data
- Storage management is about allocation and provision of resources and privileges. How do I manage the capacity to store things. We are not discussing here how I get optimal access to a piece of data on the system.
- We are interested in heterogeneous systems and domains and we are interested in access to these disparate resources with a single management system
- There are large-scale vendors who provide tools for storing and managing
- MapReduce sometimes reaches a deadlock because they run out of storage space during calculations.
- We need a layer diagram to parse storage, data, management,...

How do we know what the science communities need?

Use cases are instructive to identify what we need to address. OGF, in March is considering data/use cases to define terms.

Discussion identified that:

- MAGIC should take our findings to OGF to help further refine our understanding of storage needs.
- If I want to open a file and have 10 GB of storage available, how do I know that that amount of storage will be available?

AI: Reagan Moore will send MAGIC use case examples from the Federal community

What is the vocabulary and what do terms mean?

As well as developing a common vocabulary to provide a mutual basis for understanding and evaluation we also need to know the identity of storage elements (Are they files, vectors of Bytes, blocks, fixed block size, or objects?).

Discussion identified that:

- What does distributed mean (If I have a piece of one system and a piece of a different system, how do I put them together to address an application?)
- We are interested in both data storage on the fly and long term storage of data
- Can we identify algebras of storage management? What are the operations and how do we manage them? What do we need to support?
- What happens if a request is being rejected. We currently deal poorly with failures.

AI: Grant Miller should include the URL for MAGIC Wiki access on the MAGIC agenda.

The complete briefing may be found at:

[http://www.nitrd.gov/nitrdgroups/index.php?title=Middleware_And_Grid_Interagency_Coordination_\(MAGIC\)#title](http://www.nitrd.gov/nitrdgroups/index.php?title=Middleware_And_Grid_Interagency_Coordination_(MAGIC)#title)

Reagan Moore Data Storage Information

We work with many communities that use the iRODS data grid to implement robust, fault tolerant storage management systems. By tuning policies and procedures, the system supports a wide variety of applications:

- National Optical Astronomy Observatory – implemented a production data grid to manage images from telescopes in Chile, and archive the data in the US
- iPlant Collaborative – implemented a data grid to manage research projects and interoperate with existing repositories
- Ocean Observatory Initiative – we created a data grid to manage the archiving of sensor data streams into storage at the National Climatic Data Center
- National Climatic Data Center – using a data grid to manage ingestion of climate data records
- NASA Center for Climate Simulation – ported the MODIS data set into an iRODS data grid, and provided access to the Earth Systems Grid through a FUSE interface
- NARA Transcontinental Persistent Archive Prototype – we ported the NARA EPA collection into a distributed data management system
- Carolina Digital Repository – using iRODS to manage distribution of collections between disk and tape
- UNC-CH SILS LifeTime Library – using iRODS to support personal student digital libraries
- Texas Digital Library – using iRODS to federate multiple university libraries with a replication environment supported at TACC
- Broad Institute genomics data grid
- Wellcome Trust Sanger Institute genomics data grid
- UNC-CH genomics data grid (scale is a petabyte)
- BaBar High Energy Physics data grid – migrated 2 petabytes of data between SLAC and Lyon, France
- EUDAT replication environment
- XSEDE replication environment
- Australian Research Collaboration Service – a national data grid
- BestGrid – a national data grid in New Zealand
- French National Library storage management
- Open Science Grid storage management
- Sick Kids Hospital

- International Neuroinformatics Coordinating Facility
- Max Planck Psycholinguistics replication environment
- CyberSKA data grid (Canada)

For the prior release (Version 3.1), distributions were made to 39 countries and 38 academic institutions in the US.

We find that support for policies and procedures is as important as support for data archiving.

Other discussion:

AI: Grant Miller should include the URL for MAGIC Wiki access on the MAGIC agenda.

- MAGIC should continue to follow the survey of users discussed last month by Mine Altunay
- We should review from time to time, the MAGIC tasking from LSN

Next MAGIC Meetings

- February 6, 2:00-4:00, NSF, Room II-415
- March 6, 2:00-4:00, NSF, Room II-415