MAGIC Meeting Minutes
October 7, 2015

Attendees

Shane Canon	NERSC
Richard Carlson	DOE/SC
Susan Coughlin	ANL
Shantenu Jha	Rutgers Un.
Dan Katz	NSF
Miron Livny	OSG
Grant Miller	NCO
Tom Uram	ANL
Jack Wells	Oak Ridge
Frank Wurthwein	UCSD

Action Items

Proceedings

The meeting was chaired by Rich Carlson, DOE and Dan Katz, NSF.
Shane Canon described the computational facilities and tools at NERSC and Tom Uram described the computational facilities and tools at the Argonne Leadership Computing Facility (ALCF).

ALCF: Tom Uram

The ALCF was established in 2004 funded by DOE's Advanced Scientific Computing Research program. It operates at both Argonne and Oak Ridge. It provides 2 petascale architectures for open science. ALCF systems include:

- Mira-IBM BG/Q with 49000 nodes/786,432 cores and 786 TB RAM
- Cetus: 4096 nodes, 65,536 cores, 64 TB RAM
- Vesta: 2048 nodes, 32,768 cores, 32 TB RAM
- Cooley-Cray/NVIDIA: 126 nodes, 1512 CPU cores, 126 NVIDIA GPUs, 48 TB RAM
- Storage: 1.44 PB raw capacity

Users access the systems through peer reviewed applications:

- 60% INCITE with 5.8 billion core hours
- 30% ASCR Leadership computing challenge: DOE/SC capability computing
- 10% Directors Discretionary allocation: usually a first step to prepare for an INCITE or ASCR award

Major users of the ALCF include:

- Large Hadron Collider to simulate events to reconstruct particles produced in a collision to compare to real observables.
- Light source productivity increases through Swift data analysis to adjust the light source experiment in near-real-time
- Fireworks to define, manage, and execute scientific workflows

Challenges for the ALCF include:

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

-The LCF mission is to support large long-running jobs versus many small jobs placing requirements on the computational facilities

- Diverse computational loads
- Batch versus real-time requirements
- Data movement, archiving, curation

For the complete briefing, please see the full briefing on the MAGIC Website at: https://www.nitrd.gov/nitrdgroups/index.php?title=MAGIC_Meetings_2015 under the October 2015 Meeting

NERSC System: Shane Canon

Data are increasing rapidly and facility data often exceed local computing and networking capabilities. NERSC supports a superfacility vision where a network of connected facilities, software, and expertise enable new modes of discovery. The network supports:

- Real-time analysis and data management
- New mathematical analyses
- Fast implementations on the latest computational facilities
- Coupling data analysis and simulation
- Sharing data sets more widely
- Enabling new scientific capabilities

A common design pattern is emerging for experimental facility support. Data generation from sensors is locally processed and filtered to provide predictable data movement that enables on-the-fly calibration through analysis and modeling for real-time access and visualization supporting on-site scientists and remote users. The analysis can then be sent to storage, data management, and curation. A superfacility architecture should be deployed to apply to multiple disciplines. Challenges to this vision include:

- Unified computational architecture
- Predictable and programmable networks
- Workflows for seamless data movement
- Productive user environments for data analysis

Cori provides a unified system for big data and HPC. It provides:

- Data partition with Xeon Processors and larger memory (128Gb)
- HPC partition with Xeon Phi (KNL) Processors
- Common bandwidth interconnection
- Common access to NVRAM Burst Buffer and high-bandwidth parallel file system
- Ability to support custom user-defined images
- Advanced gateways

Such unified systems enable the coupling of experimental data with simulation and modeling.

Predictable networking is facilitated by use of the science DMZ architecture and the use of Software Defined Networking (SDN). High-bandwidth networking (100G+) enables faster data movement from an experiment to computation. New techniques are developed for moving and managing data through the complex memory and storage hierarchy.

The software ecosystem on the HPC platforms is optimized for large scale simulation and modeling with integration of math libraries, debuggers, visualization tools, compilers, parallel frameworks, parallel I/O libraries and performance/profiling tools. Shifter is a container for HPC that enables users to bring a customized OS environment and software stack to an HPC system. It is used for High Energy Physics, Cosmology and bioinformatics and light source applications. Shifter has been observed to significantly increase performance of applications.

Conclusions:
NERSC has seen a new focus on time-sensitive applications. Science engagement, meaningful interaction with scientists will become increasingly important. The concept of users will increase to include other computational facilities. Then, coordinated user support across facilities will need to be scaled and responsibilities defined. Orchestration across facilities will need to be coordinated at multiple levels from resources to outages.

For the complete briefing, please see:
https://www.nitrd.gov/nitrdgroups/index.php?title=MAGIC_Meetings_2015 under the October 2015 Meeting

Potential MAGIC tasking form the LSN

MAGIC developed a number of topics they would like LSN to task them with for the upcoming year. These topics include:

- OSG growth and evolution; XSEDE components: Wrangler (data intensive computing), COMET infrastructure, JETSTREAM Infrastructure
- Advanced computational modeling and cloud computing: Radically new techniques and standards. Invite involved organizations to talk.
- Federation of repositories across federal agencies. Current applications of iRODS include NASA, NOAA, NOAO, NSF projects. Keith Marzullo potential speaker
- E2E encryption in the cloud
- Identity management and access control in virtual environments and clouds

Meetings:
October 15-16: Software management workshop, Washington DC

Next MAGIC Meeting
November 18, 1:30-3:30, SC15, Hilton Hotel, Austin, Room 414