MAGIC Meeting Minutes
July 5, 2017

Attendees

Gonzalo Rodrigo	IBL
Katie Antypas	NERSC
Rich Carlson	DOE/SC
Dan Gunter	IBL
Dave Goodwin	DOE/SC
Shantenu Jha	Rutgers Univ.
Jim Kirby	NAVY
David Martin	Argonne
Grant Miller	NCO
Jerry Perez	TTU
Rajiv Ramnath	NSF
Eric Rees	TTU
Don Riley	UMD
Alan Sill	TTU
Kevin Thompson	NSF
Jack Wells	ORNL

Action Items
Proceedings
This MAGIC meeting was chaired by Rich Carlson (DOE/SC) and Rajiv Ramnath (NSF).

Presentation: Katie Antypas (NERSC), Data Department Head (slide presentation available)
Introduction to NERSC's Data Department and its efforts to support data intensive workloads
NERSC runs HPC resources and provides expertise and advice for users. NERSC recently formed its Data Department to focus on data intensive science and has deployed a new data initiative to enable large-scale, data-intensive science discoveries.

A long-time supporter of data intensive science, NERSC has recently supported users with data intensive projects (e.g., ALS light source). Historically, 2 systems existed: 1) a compute intensive system and 2) a smaller cluster for data intensive science. Today, the data being generated is too large for current facilities and many wish to combine large scale simulation and data analysis. There is a need to incorporate data from experiments as well as simulations.

NERSC Plan- Combine capabilities into 1 system
CORI, NERSC's latest system, has capabilities that support data intensive science. NERSC is transforming how the traditional HPC system supports data intensive science by:

- streaming data directly into super computers at a much higher bandwidth by replacing RSIP technology on Network nodes
    - Burst buffer architecture: place flash on I/O node, not compute node
    - Burst buffer is gaining momentum and particularly impacts the speed of read-intensive applications
    - Burst buffer can reserve space on system

- providing a layer of nonvolatile storage/flash, which is now in the super computer. CORI Burst buffer system (1.5 PB, 1.5 TBsec) increases available I/O bandwidth. More cost effective to buy bandwidth with flash or nonvolatile technology.
- providing virtualization capabilities that allow user to bring own software stack to the super computer
- providing more login nodes for managing advanced workflows
- supporting real-time queues - 32 nodes reserved for applicants

NERSC User Support

- NERSC Exascale Science Application Program (NESAP) for Data prepares teams to run on newer architecture with high bandwidth memory (KNL). Some of the teams are already using KNL for data intensive simulations (e.g., Deep Learning at 15PF; Celeste and Galactos).
- SPIN Platform: Edge services to support data intensive science that's running on supercomputers. Due to increased demand for these services, NERSC is providing formal infrastructure support. Users control the applications (launch gateways, WF, databases based on "docker" containers)
- NERSC -9 Project delivery in 2020 – released RFP requiring system that supports extreme data science and computing

Future Data Requirements

Scientists would like:

- More and improved support for analysis tools as analysis code doesn't always scale to the HPC systems. Many analysis tools significantly differ from traditional HPC simulation.
- New ways to analyze data; experiencing huge data surges and need to incorporate new approaches like advanced stats or machine learning. Some are building their own.
- Work flow accommodation: Work flow in simulation and analysis are more complex and need accommodation on HPC systems (eg. work flow sometimes starts outside HPC system and needs to be transferred over WAN. Or need databases or gateways to manage data.)
- HPC at experimental facilities to guide experiments in real-time, so need co-scheduling between DOE facilities.

Summary

- Research impeded by authentication/trust/ID management between experimental facilities and NERSC
- Analytics software needs to be scalable
- Need more common tools and more robust system so scientists can run analyses; Need co-scheduling for different resources (burst buffer, storage, bandwidth)

Discussion

- Goal: still unified architecture because of the difficulty in separating workloads. This trend is supported by scientists who want to do analysis on same processors that do data analysis and simulation.
- Open question: how easily data intensive workloads will run on architectures like KNL.
- Issues of cost, sustainability, data transfer options
- Where does NERSC fit into larger ecosystem of commercial cloud, data and compute services? How can infrastructure be supported going forward?

Potential MAGIC Tasking FY 2018

Group discussion and modifications are reflected in the list below. **Newly added topics are in red.**

- Existing/developing virtual environments: OSG, OGF, GENI, FutureGrid, Internet2Net + environment, supercomputer environments
- Convene the OSG, CERN, OGF… communities to discuss their different approaches and what has worked/what has not worked.
- Bring the NSF funded cloud environments into the MAGIC discussions to represent academic community interests
- University community researchers and providers to identify current capabilities and desired future capabilities.
- SDN developers to identify how their developing technology might impact virtual environments and distributed resources/distributed processing
- Cooperation with commercial cloud providers
- Data movement and data management. Middleware is expediting movement of data across collaborating groups and among science disciplines. An example is cooperation among NSF data hubs for moving data to/from supercomputer centers. CASC is participating in this effort.
- Evolving Identity Management (IdM)
- **Improving the reliability of middleware and grid environments. Software and networking are critical components for improving reliability.**
- **Data intensive science: impact on distributive computing and large scale supercomputing facilities. How is this environment changing (possible workshop)**

Discussion topics

- Data movement and data management. Middleware is expediting movement of data across collaborating groups and among science disciplines. An example is cooperation among NSF data hubs for moving data to/from supercomputer centers. One member noted that CASC is participating in this effort.
- Evolving Identity Management (IdM)
- **Explore more efficient operation of data centers for universities and labs by incorporating up –to-date knowledge.**
    - Extracting and building better communication between cloud research projects and data center operation. (e.g., CloudLab and Chameleon and counter examples: Shifter, Singularity).
- **Identify how commercial resources (e.g., cloud environments) can be used/integrated into science environments, including involving experts from commercial cloud providers**
- **How NSF compute facilities are addressing data intensive science community**
    - **Labs are looking into building more comprehensive compute environment that crosses physical domain boundaries of each lab**

MAGIC Roundtable

DOE: Rich Carlson

No news to report.

Argonne: David Martin

MPI hosted tutorials at the Argonne TCS Center

TTU: Allan Sill

IEEE magazine wants to expand community involvement/ volunteers to contribute to website. It is requesting ideas for how to increase volunteer engagement from the IEEE cloud community.

Meetings of Interest

July 9 – 13	Practice & Experience in Advanced Research Computing, New Orleans, LA. Rajiv will try to report back to the group.
July 30-August 11	Argonne Training Program for Extreme-Scale Computing, St. Charles, IL

Next MAGIC Meeting

August 2, 12:00-2:00 p.m. EDT, NSF, Room TBD