***The government seeks individual input; attendees/participants may provide individual advice only*.**

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes[1]
February 6, 2019, 12-2 pm
NCO, 490 L'Enfant Plaza, Ste. 8001
Washington, D.C. 20024

Participants (*In-Person Participants)

Deb Agarwal (LBNL)	Joyce Lee (NCO)*
Lisa Arafune (CASC)	Brian Lin (UW-Madison)
Wes Bethel (LBNL)	Miron Livny (UW-Madison)
Richard Carlson (DOE/SC)	David Martin (ANL)
Dhruva Chakravorty (TAMU)	Gilberto Pastorello (LBNL)
Vipin Chaudhary (NSF)	Don Petravick (NCSA)
Jill Gemmill (Clemson)	Ryan Prout (ORNL)
Sharon Broude Geva (UMich)	Lavanya Ramkrishnan (LBNL)
Alexander Hexemer (LBNL)	Hakizumwami Birali Runesha (UChicago)
Florence Hudson (FDHint)	Sonia Sachs (DOE/SC)
Shantenu Jha (BNL)	Suhas Somnath (ORNL)
Margaret Johnson (NCSA)	Von Welch (IU)
Hari Krishan (LBNL)	

Proceedings
This meeting was chaired by Richard Carlson (DOE/SC) and Vipin Chaudhary (NSF). December 5, 2018 meeting minutes were approved.

Speaker Series: Data Life Cycle
- *Deduce: Distributed Dynamic Data Analytics infrastructure for Collaborative Environments*- Deb Agarwal (PI), Department Head and Senior Scientist and Lavanya Ramakrishan, Staff Scientist, Lawrence Berkeley National Laboratory[2]
- *Scaling Data-Driven Scientific Discovery – A Data Lifecycle View*- Suhas Somnath, Computer Scientist, Advanced Data and Workflows Group, National Center for Computational Sciences, Oak Ridge National Laboratory

[1] Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.
[2] Team also includes Bhavna Arora, Suren Byna, Stephen Bailey, Boris Faybishenko, Abdelrahman Elabashandy, Devarshi Ghoshal, Juliane Mueller, Gilberto Pastorello, Talita Perciano, Tom Powell, Alex Romosan, Craig Tull and John Wu.

- *Addressing the Data Challenges at the Advanced Light Source*- Alexander Hexemer, Senior Staff Scientist, Program Lead for Computing, Advanced Light Source, Lawrence Berkeley National Laboratory

Deduce: Distributed Dynamic Data Analytics Infrastructure for Collaborative Environments **- Deb Agarwal, Lavanya Ramakrishnan**

How we deal with the data life cycle: Examples of engagements (end-to-end earth science):

- AmeriFlux and FLUXNET Carbon Flux: brings together data collected by 800 towers globally; working with them to organize data end-to-end from volunteer collaborators.
- NGEE Tropics: understanding climate change variables and vegetation's reactions to climate change.
- Understanding watershed function as a whole: Complex, multi-domain, multi-scale problem: Bringing together many disciplines (e.g., meteorology, vegetation, geophysics, hydrology) and how nutrients cycle through the watershed (include contaminant variables, microbes).
- Many projects with difficult characteristics such as many data contributors, large scale data, domain issues, scale issues.
- Almost always trying to connect the end-to-end system where we are bringing together models, data processing, storage and archiving. Trying to bring data together: Involves internal and external data and working with scientists.

View of Data Lifecycle

Dynamic system: e.g., back and forth process (analysis finds issues with data, so do further collection and processing).

Projects (In addition to projects involving end-to-end problems):

- Projects that address different aspects of data lifecycle (how we manage, visualize, build work flows, connect up to modeling and simulation, how use data to build parameters and surrogates for models).
- Data archive run for ESS-DIVE (DOE's Environmental Systems Science Data Infrastructure for a Virtual Ecosystem).
- Usable data abstractions.
- Science capsules: work with users to understand use cases and build out data life cycle capabilities based on users' needs.

Usability:

- Try to understand use cases, how use data, needed tools and building around it (e.g., TiGrESS, Deduce projects).
- Need to build uniform integrated views on data: come up with ways to integrate external and internal data into a single view for scientists to ask science-driven questions.

Challenges of Data Life Cycle and hidden problems.

E.g., Fluxnet: built 2 major global data sets with Europeans: LaThuile 2007 and FLUXNET 2015

- Lessons learned: underlying data is continually changing in small ways.
    - E.g., Set of Correlation plots for 1 variable, just looking at site years where it was same for that variable across 2 datasets.
- Why change?
    - AQ/AC or processing on data that is continually re-visited throughout cycle.

    - Sometimes better science understanding of raw data, so regenerate data.
- Dimensions of change. Factors:
    - Values: if archived, we don't know that it has changed
    - Organization of data holding system: file system and metadata changes
    - Evolution of data in real-time
    - Impact: lack good tools for assessing impact

Research

- Deduce:
    - How understand the impact of the change
    - Track and measure data change at scale
    - Ways data archives and other data holders represent change to data users
    - Manage data change workflows at scale on HPC machines
    - Example: impact of change on model that is running using that data as an input parameter
        - Understanding the above-ground biomass of different species of trees and their reaction
        - 2 different versions of meteorological drivers received from the data source: significant impact on different species. The tail of small change at beginning of data could have a major change throughout the data cycle.
- One way to address changing data: Look at how much information we need to derive a quantity or come up with an answer? Use information theory to understand how to figure it out. If sufficient data, perhaps can be more resilient to the changes.
- Dac-Man framework: built framework to inspect data and identify changes. Building out metadata representation to enable data sources and users to inspect data set and quantify changes. Also building tools to assess impact.
- User Perceptions of Statistical Change Analyses:
    - Try to understand how to represent change to the users to enable their understanding.
- Helps to see data life cycle challenges in the context of individual problem as it highlights the many types of problems in this system:
    - User Interface Mockups and worked with users: ways to represent change so users can understand and take action on findings
- Change representation software on Github (DacMan: https://github.com/dgoshal-lbl/dac-man)

***Scaling Data-driven Scientific Discovery in the Big-Data Era – A Data Lifecycle View* - Suhas Somnath-**
Paradigms of scientific discovery

1) Simulation, 2) observation experiments, and 3) data analysis, which has begun transforming the landscape for numerous scientific disciplines and relies upon large volumes of well curated data. There has been an explosion in data volumes, along with faster supercomputers for larger simulations.

Current Data Life Cycle in Experimental Facilities

- Challenge: Data lagging (post-processing and managing data), especially small facilities which are generally driven by commercially available instruments and researchers use a variety of software. After analysis and pre-processing, more data files are stored and forgotten on hard drives.

Acquisition

Big Data and Information Transfer in Microscopy:

- Standard Mode: easier to store "data trickle" on hard drive; could lead to incorrect views
- Full Information Acquisition: store entire data stream in storage and use variety of data analytic tools (e.g., ML) to see larger trends/correlations in data (General Mode, "G Mode")
    - Can extend to other areas of observational sciences
    - Important to keep all data: Extracting trends in diffraction patterns in Scanning Transmission Electron Microscopy (STEM) data. Can understand more from the data.

Standardization:

Multitude of File Formats and Data Types

- Used proprietary software, so incompatible. Not suitable for sharing, archiving.
- Need to agree upon data model (abstract representation of data) and file format (container storing data). Can reuse modalities. Brings community together to build ecosystem of service for sharing, publishing, etc. While many communities have done so, it is still a challenge for others.
- Example: Universal Spectroscopic and Imaging Data (USID)
    - Flatten original N-dimensional form into 2-dimension form (can reuse same techniques across different communities)
    - Hierarchical Data Format (HDF5): "file system within file" that stores metadata.
        - Many communities (e.g., light sources) use HDF5 to store data.
        - Open, free, easily accessible

Analysis

Challenges with Scientific Software

- Proprietary software is closed-source, antiquated, and expensive to install.
- Users: not have access to software. Difficulty using scripting/programming, no laptop access.

Solution: 2-program software strategy

1) Exploratory: driven by domain scientists who write their own software; run simpler but approximate algorithms. Research-driven arm of software (e.g., Python)
2) Production: run on HPC or in computing cluster. Instrument-to-HPC workflows (e.g., Beam, Xi-CAM) connect instrument to computing tier. Facilitates experiments and analyses.

Towards Open and Reproducible Science

- Software and data are insufficient
    - Need tools (e.g. Jupyter notebook– when applied with data sets and manuscript, goes a long way towards reproducibility. Adding software container helps even more).

Transfer & Storage

- Need to store in centralized repository, failure-resistant, close to compute resources. Anything that doesn't need to be used or is rarely used can be archived.
- Data transfer to Centralized Repository
    - Labs have different combination of instruments. Researchers want instruments kept offline, but want to stream data to an external, centralized storage repository.
        - Solution: Instruments stay offline but connected to a bridge device that transmits data through dedicated fiber line. Then have access to different computing resources to move forward in data life cycle.

Scalable Data Efforts at ORNL:

Collection and Management: iTunes for Scientific Data- to organize data and search via metadata and access data from any computer. One Solution: DataFed which has many capabilities:

- Organize data sets based on project. Each data sets associated with rich meta data, which can be used to search, for example, a data set generated by specific person.
- Software is capable of fine-grained access control and sharing for project, etc.
- Federated identity via Globus.

Publishing and Cataloguing: Making finished product available to community
Prototype of Data Publishing:

- Bring data files and submit to web application.
- Final stage: review process- same as journal review process and also check for metadata, open, community-defined standard and that data is correct and of good quality.
- Cataloguing: See Prototype of ORNL Data Catalog. Show that it is acceptable and can be mined (use catalogue). Can be used to understand broader trends in data.

Data Mining – last step of data life cycle
Scalable ML: Since large data sets, need scalable version of packages or new packages written to address large data sets (e.g., pbdR, Non-negative Matrix Factorization) capable of scalable ML.
Scalable Deep Learning: Can scale up PyTorch networks by using Horovod (distributed Deep Learning Framework) or MENNDL(Multi-Node Evolutionary Neural Networks for Deep Learning).

***Addressing the Data Challenges at the Advanced Light Source (ALS)*- Alexander Hexemer**
Theme: We can't do it alone, we are putting together a framework together across multiple facilities. Working with LBNL, ANL, SLAC, BNL

Center for Advanced Mathematics for Energy Research Applications (CAMERA) project

- Bring latest mathematic theory to domain scientists and implement. Run on multi CPU/GPU, make everything open source and try to collaborate across facilities.
- Autonomous data collection collaboration BNL letting computers steer experiment. Involves ML, typography, imaging analysis.
- Important due to challenges for user facilities:
    - Wish to provide quick feedback, provide data from different modalities and experiments, build flexible workflows and implement latest math and algorithms, access across facilities.
    - 20% of users are new and many lack HPC experience.
    - Complex raw data (prediction ALS)
        - Jumps from CY 2018-2024: couple PB/year to CY2028: 30 PB/year because of new diffraction limited light sources coming to ALS.
        - Old ALS-U, then upgrades (3 orders of magnitude of brightness (x-rays per unit area).
        - APS in Chicago is facing the same challenge of going to APS to APS-U APS.

Light Source Data Working Group:
Formed to develop standards for data storage and file formats. Try to do real-time computing. Create reports of how much data will produce, computing requirements, etc.

Local collaborations: Development of the Data Movement in Collaboration with other LBNL divisions Supercomputer centers, Lab science IT, ESNet, Data Acquisition Group, CAMERA

ALS: Collects time resolved time diffraction patterns. Want model to extract science from raw data.
<u>Full data lifecycle currently in development at ALS</u>

- Detector running at very high speeds. Perform pre-processing likely close to beam line (depending on data). Stream some of data using Globus onto compute cluster which is locally in ALS. Will have some local storage, but mostly temporary.
- Also do it to collect metadata from beamline. Package in event based model (BNL's Data Broker) and ship it.
- Once leave this environment, stream data to a compute cluster. Not just move data or use to launch hard drives.
- Graphical user interface (Xi-Cam) implemented to deploy workflows and do data analysis; not move data away from compute cluster. Plug-in framework to provide scientists with plug-in capable system to write own code, but we take care of data work, integration of HPC computing.
    - Community-maintainable platform for new analysis, deployment of mathematical algorithms for CAMERA. Recently developed variety of plug-ins (e.g., tomography)
    - Other beam lines are starting to use Xi-Cam.
        - E.g., work with SSRL –scan crystals with x-ray beam and convert in real-time x-ray data into correct space. Xi-Cam provides crystal structure in real-time, so know immediately if experiment works (or not). Saves valuable beam time.
    - Remote Execution: GISAXS workshop in Bayreuth in collaboration with CAMERA
        - Simulation: trying to simplify access to HPC. Want to connect to and do calculations at a computer in Switzerland using Xi-Cam for remote control.
    - How interact with more complex systems:
        - High-throughput NEXAFS workflow collaboration with Materials Project.
            - Can stream measurements directly from Xi-Cam into the analysis framework and into the materials project. Connect input, workflows and analysis through the Xi-Cam link workflow.
    - <u>Deep learning for X-ray Scattering</u>: Provide more guidance for users to make right choices immediately. Must treat data differently; using different methods and analysis. Train neural networks which deploy in Xi-Cam; classify as soon as data comes in.
        - Why? 20% users are new and need help and increased data rates with new data sources coming in necessitates faster decisions on sorting and whether to keep data.
- <u>Use Case in Development</u>: build new x-ray photon spectroscopy beamline. Collaborative approach to develop beamline and working on this data.
    - Interface: Xi-CAM to run beamline and deploy workflows on remote computing.
    - Control: working with BNL to implement Bluesky and DataBroker
    - HPC Code for XPCS (ANL)

<u>Vision</u>: Make science discoveries as easy as possible (not have to teach everyone how to explore science)

- Implement visual and Interactive systems
- Do multi-modal so have different spectroscopy aspects in data
- Implement physics and math engines by design
- Full integration of ML

Discussion

Top 2 challenges to realizing your vision?

- Overall challenge: having all needed tools to go from end-to-end. Moving from current situation to where it's reasonably routine. Trying to put together a data pipeline from instrument to user service to analytics, ML, etc. is very challenging and takes much expertise.
- Who will pay for holding visitors' data in data catalogue.
- How to train domain scientists to view each other's data sets.
- Motivating younger scientists to have more exposure to computing; perhaps curriculum changes are needed.
- Commercial proprietary data formats in software. Force vendors to provide more open data formats; would benefit scientists nationwide. Need support of from funding agencies and journals to require community standard.
- Who will pay for data storage? Note: Future increase of compute requirements
- Funding agencies need to mandate that awardees have a data plan, not just in hard drives. Also have data stored in archive and funding necessary components. Helps shift community.

Cultural Change

- Cultural change from individual data sets to moving data to compute facility and build data into major infrastructures (working in the community). Need to start prepping now for future change (e.g., more light sources).
    - Also ensure data credit, data publishing work well when no longer in complete control of data. How would it positively impact scholarly reputation.
- ALS: How to correctly throw away raw data.

Data Life Cycle Planning

Identify gaps in coverage and potential speakers

- February: Framing the issue. Talked to data sources and providers of instruments
- March: Science community to discuss use cases
- April on: need to identify topics

Future Session Topics (From Outline Summary)

- Triaging Data: determining what data to throw away, archive and save. How will this work? Challenges in this area?
- Data Analysis and cross facility users: Useful to understand how science communities will be able to reuse data. How make data from someone else's experiment available to another scientist to analyze without re-running data collection process. Is this something that communities thinking of (facilities and science communities could address)
- Data Reproducibility: always a concern
- Data Storage: cost
- Provenance: Which algorithms have been used to transform data? HDF5 also has more detailed provenance information or does workflow management system provide? Other elements?
- Privacy and Security: Have not addressed yet. If sharing data, need to give rights to certain colleagues and prevent access to others? What about data you have collected and stored outside of your facilities?

Potential topics and speakers

April: Data use/re-use and data provenance/integrity

- Jane Greenberg (Drexel Metadata Research Center) who has a number of grants
- Margaret Johnson (NCSA) and Don Petravick (Data "tagging"- understand more aspects of data as use it; related to data provenance and integrity)
- Amadeo Perauvio (Free Electron Laser): which data to throw away (Alex Hexemer referral)
- Data provenance speaker

Look at summary list offline to identify potential set of topics. Let us know of any topic that has been missed.

MAGIC Tasking (CY19)

Report on new networking capabilities focusing on containerization and DevOps. MAGIC reports to NITRD's Large Scale Networking IWG. Reports will be derived from the MAGIC minutes and presentations that are posted on the MAGIC weblink. Once draft is written, will circulate it to MAGIC team for review.

- Containerization Report: Dhruva Chakravorty (TAMU) has volunteered to compose a short report on containerization.
- DevOps Report: We are looking for a volunteer to put together the DevOps report.

Roundtable

- Agencies are catching up and will likely have announcements in the next meeting.
- Academic Roundtable/Curriculum changes
  - MAGIC is engaging with the academic community for the MAGIC community's awareness and it is an issue related to workforce development
  - Dhruva Chakravorty (TAMU): In high-performance research computing unit. Annually, have 5-6k students across Texas working on WebEx platform and doing online labs. Facing challenges translating TAMU's efforts into curriculum. Would be interested in hearing from others facing the same challenges.

Next meeting

March 6 (12 noon EST)