



Lawrence Berkeley  
National Laboratory



# HPCOR 2014

Data Management Policies  
Co-Chairs: Bill Allcock, Julia White

# Contributors

- Bill Allcock, ANL
- Sasha Ames, LLNL
- Ashley Barker, ORNL
- Laura Biven, DOE
- Jeff Broughton, LBL
- Bill Collins, SNL
- Mark Gary, LLNL
- Kyle Lamb, LANL
- Dee Magnoni, LANL
- Joel Stevenson, SNL
- Julia White, ORNL



# Top Discussion Points

- Currently, data management at the facilities is largely capacity and I/O performance planning.
- It is more than capacity, it is also throughput and metadata performance.
- Data management is more than just the items above, but includes things such as lifetime, provenance, curation and access.
- As we expand the definition of data management we will need to define additional policies.
- What capabilities and resources should the facilities provide for policies that will exist and what are the resulting policies at the facility level?
- Facilities role might include education, enablement, collaboration between facilities on possible solutions or enabling tools.
- We really don't know what this is going to look like today.
- Scheduling of data access; Making data equal to computation

# Challenges and Gaps

- Data import / export capabilities (bringing us disks or tapes or filesystems)
- There are more users, with more data, who want to use it in different ways than in the past. Adapting from “lots of data” to “big data”
- The size of data is growing but so are the access models and the services that are requested on the data.
- The larger data management problem of curation and provenance can be enabled and assisted by the facilities, but the users must be the ultimate driver since they have the domain expertise.
- Risk analysis of saving vs. regenerating the data, whether or not to save the data at all, etc..
- Dearth of storage expertise and competition with industry

# Best Practices

- Make sure our current and any future data management policies, services, resources are explicit and clear.
- Asking up front what the users projected data needs will be and make that a first class citizen in planning.
- Educate the top resource users on the cost/value of the resources they are consuming to encourage wise usage of the resource; Store only what is needed; Improve I/O performance;
- Embedding staff into the science teams.

# Opportunities and Collaborations

- Facilities role might include education, enablement, collaboration between facilities on possible solutions or enabling tools.
- Make data management plans accessible to the facilities for the purposes of resource planning.
- Request metadata and text files providing research context.
- Work with the external facilities that are large producers of data (light sources).
- Collaborate with work flow developers and researchers to incorporate and standardize metadata for long term storage.

## What are your major strategies and initiatives over the next 5 & 10 years? How do they affect staffing levels?

- Adding two new optional layers to our data hierarchy
- Burstbuffer layer between main memory and parallel file system
- Campaign storage as a disk based storage that allows for quicker recall of data that has been moved off from parallel file systems
- Efforts shall be made to allow in situ analysis of data on the burst buffers
- Efforts shall be made to allow data reduction and analysis in campaign storage
- Building larger infrastructure to accomodate high bandwidth data movement
- Staffing has to be augmented to support the extra data layers while still supporting existing data layers.

## What are your major strategies and initiatives over the next 5 & 10 years? How do they affect staffing levels?

- Convene a Data Working Group comprised of personnel from Research Library, HPC, ADIT, Big Science, Small Science and other interested parties
- Create a data policy
- Plan and create Researcher toolkit
- Increased need for data scientists and programmers

## What are your major strategies and initiatives over the next 5 & 10 years? How do they affect staffing levels?

- Big Data efforts - we are trying to find where big data can leverage our “lots of data” expertise.
- Integrating storage as a player with center-wide HPC scheduling.
- Staffing levels will likely remain constant.

# What are your current efforts and/or site configuration in this area?

- We have a burst buffer test bed that is in use for fast forward testing
- Initial implementation of campaign storage in production on small scale using commodity disk and RAID 6
- High bandwidth data movement project testbed
- Seeking approval for DMPTool
- Seeking approval for a sync and share solution for collaboration
- Gathering together data policies from existing institutions to pull best practices
- Tracking DOE SC policy development

# What are your current efforts and/or site configuration in this area?

- We are working on non-impactful query mechanisms to increase speed and minimize load on file system metadata.
- We are adding crash barriers to our existing soft archive quotas
- We are adding situational intelligence to our data movement tools

# What are your mandates and constraints? (costs vs. retention vs. access)

- Prior mandates were enforced by performance, looking to utilize lower cost hardware coupled with erasure code for reliability and commodities of scale to enable higher bandwidth at lower overall cost
- Quotas on archive/Campaign storage are likely to be needed.
- Growing file system sizes to meet performance may allow for longer purge policies, and purge policies are changing
- OSTP mandate on federally funded research
- Lack time, staff, resources, current policy and roadmap [for data]

# What are your mandates and constraints? (costs vs. retention vs. access)

- Mandate - keep archive data forever. Capacity is relatively cheap so this isn't a huge issue - it is bandwidth that is expensive
- On our file systems we have 60 days since last access purge policies. Our real constraint continues to be how quickly we can purge and manipulate files (file system metadata performance)

# How to do you forecast future needs and requirements?

- User input via interactive meetings
- Technology trend analysis
- Input from program offices
- The Data Working Group will be charged with outlining current and future requirements
- From an HPC perspective we relate most things to platform memory. Burst Buffer deployments may well turn these metrics on their head.
- Tactically we ask for information on data requirements (including retention and movement) during CFP processes.

# What **opportunities** exist for productive collaborations among DOE HPC centers?

- High bandwidth data movement tools and standards
- Enabling of larger data sharing capabilities
- Sharing of technological ideas and pitfalls

## What are the biggest **challenges** and gaps between what you can do today and what will be required in 5 - 10 years?

- Performance of infrastructure is not scaling to meet the needs of cluster designs
- Shift toward large memory systems has a major effect on requirements of data infrastructure