



Lawrence Berkeley
National Laboratory



HPCCOR 2014

Infrastructure (D2SA)

Co-Chairs:

Christopher Beggio, Sandia National Laboratories

Robin Goldstone, Lawrence Livermore National Laboratories

Contributors

- **Bill Allcock**, Argonne Leadership Computing Facility
- **Chris Beggio**, Sandia National Laboratories
- **Clay England**, Oak Ridge Leadership Computing Facility
- **Doug Fuller**, Oak Ridge Leadership Computing Facility
- **Robin Goldstone**, Lawrence Livermore National Laboratory
- **Jason Hick**, National Energy Research Scientific Computing Center
- **Kyle Lamb**, Los Alamos National Laboratory
- **Jim Silva**, Lawrence Livermore National Laboratory
- **Jay Srinivasan**, National Energy Research Scientific Computing Center
- **Ilana Stern**, National Center for Atmospheric Research
- **Craig Ulmer**, Sandia National Laboratories
- **Venkat Vishwanath**, Argonne Leadership Computing Facility

1. How does data-driven science impact your HPC network and storage architectures (incl. archive)? Are significant changes required or do you just need "bigger" pipes and "more" storage?

- HPC facilities have good handle on how to ingest/write data but need to focus more on read/query issues.
- Need better middleware layer to facilitate data movement between storage layers. This is not just HSM. Need locality-awareness, scheduling integration, etc. Get the data where it is needed, when it is needed.
- No one-size-fits-all solution. Each site has unique application/user driven requirements that are influencing their approaches.
- Might need custom solutions initially but all labs would prefer to leverage general/common solutions for HPC and data science where appropriate.

2. How is data-driven science impacting your HPC planning processes? Have you formally integrated data science requirements into your planning process or are these efforts currently being pursued "on the side?"

- Most sites' initial efforts were being done on the side i.e. with repurposed hardware, driven in part by funding and lack of formal requirements.
- Increasing need is being recognized and now many sites are formally planning for and in some cases procuring systems targeted at data science.
- Would prefer to leverage existing infrastructure and systems wherever possible. Want an integrated environment but may need some unique components to make it work.
- Networking requirements are still being addressed after the fact or not at all.

3. How is data-driven science affecting your HPC procurements? Is the spending balance being shifted more toward storage and alternative data-centric architectures? Is more money needed/available to meet these evolving requirements?

- Some sites buying platforms targeted at data science, others are focused on leveraging existing HPC systems in a more data-centric way.
- Viz systems may be good initial candidate to evolve into more general purpose data analytics systems due to sharing common hardware attributes (more memory, better I/O). Emerging data science requirements may drive better “traditional” viz systems in the future.
- Most sites see the need to integrate SSDs into HPC systems for burst buffers which will also benefit data science. However, this will be at the expense of other HPC system components unless more money is provided.

4. What alternative storage technologies and file systems are being considered to support big data storage requirements?

- All sites are investigating options in this area, including the following.
- Burst Buffer solutions
 - Node-local SSD
 - PCI-e
 - IBM Ultra DIMM
 - SSD in other parts of cluster e.g. I/O Nodes
 - Future NVM technologies e.g. PCM
- EMC Viper
- High density node-local storage e.g. HP SL4500
- JBOD to replace RBOD
- ZFS
 - Lustre, HDFS, Gluster on top of ZFS
 - Storage appliance w/ embedded processing eg Xyratec
- Ceph
- Gluster

5. How are WAN providers responding to inter-facility large scale data movement requirements?

- Need on-demand network provisioning, integrated with compute and storage resource scheduling
- Last-mile connectivity: BW, reliability
- Edge security devices required in some cases but can inhibit performance. Can't just blame the firewall.
- Ideally sites can isolate HPC/data science network from enterprise to meet requirements: bandwidth, security model, etc.
- Planning for network needs is typically done separately from platform planning, often as an after-thought.

6. Do data-centric workflows impact your security posture, e.g. need for firewall bypass to efficiently transfer large data flows across security boundaries, real-time data acquisition, etc.

- Science DMZ model is working for many sites.
- Need federated authentication/authorization solutions.
- Security model: some sites do not want to expose production resources directly to Internet, “data DMZ” may be necessary, but results in storage duplication and reduced end-to-end throughput.

7. Are commercial cloud and big data solutions applicable to the scientific community or do we need to "roll our own?"

- Private cloud model is being used effectively by some sites.
- Public cloud raises concerns about bypassing data review and releases processes for scientific data.
 - Public cloud solutions effective for enterprise services.
- Many commercial solutions do not target large scale scientific computing and fall short, typically in ability to accommodate our capacity and performance requirements.
- Lack of control over external providers can impact:
 - Availability
 - Performance
 - Security

Top Findings

Opportunities

- All sites exploring the same topics, great opportunity for info exchange and collaboration.
- Leverage solutions from other sites: we can't afford to reinvent the wheel.
- These early investigations will position us to meet emerging customer requirements.

Best Practices

- Leverage HPC solutions wherever practical.
- Apply extensive storage experience and talent to address data science challenges; we know a lot about data, this is just another kind of data we need to understand.

Challenges

- Wide diversity of data science applications each with unique requirements.
- Data science savvy staff are targets for Google.
- No direct funding for data science investigations.