Scientific Computing for FY2020-23

Gaps in Capabilities

or,

Opportunities for Science through Initiatives in Computing

Chip Watson

Head of Scientific Computing

Nov 27-28, 2018

Missing Capabilities and Expertise

1. AI and Deep Neural Networks

- Accelerator modeling and tuning
- Detector monitoring
- Event reconstruction
- Detector simulation

We have adequate hardware resources to get started (end of life LQCD GPU nodes or idle time on farm nodes from time to time), but don't yet have the staff and expertise (options: hire, or re-prioritize and train).

Identifying a likely high impact project (more uptime, data per funded week of accelerator operations) and building a collaboration between IT division and one of the other divisions would be a good way to start.

Additional Software Support for Physics

The LQCD SciDAC and Exascale projects support 2.8 FTE of application support within Scientific Computing. Support for ENP is funding primarily for workflow tools for the batch system and for data movement to/from tape and to/from NERSC.

Expanded support for all 4 halls could yield significant benefits.

2. Add 1 FTE to support additional tools for data analysis and other work currently handled by fractions of physicists
   - Generic tools like NumPy, Juniper Notebooks, in-memory databases, debuggers and profilers, …
   - More specific components useful to all halls such as more powerful and general purpose data catalogs for large datasets
   - Community tools like OSG tools, GEANT, CERN software, …

Enhanced Hardware Capability

3. Provision for higher throughput interactive computing

Aim: enable experimentalists to be more productive.

Costs:

- Additional disk space (so data sets stay resident)
- Possibly SSD for faster skims
- Dedicated powerful interactive nodes, or use of pre-emption to quickly shift capacity to interactive use when appropriate
- (additional complexity would put more demands on a lean staff, but this could potentially be grown slowly to minimize a sudden large increase in trouble tickets)

New Capability for the Big Experiments

4. Near Online Computing

(has been discussed internally between CLAS-12 and IT)

Building a large 3rd level trigger cluster in a counting house could allow running at higher beam intensity or trigger rates, (if the detectors could support it) but would lead to an underutilized resource when beam is off.

Dynamically carving a level-3-like resource out of the farm would enable higher overall utilization, and allow the size to be adjusted to the current running conditions. Nodes are already supported.

Hardware would be simple, but staffing might be a challenge. We are a lean organization (think 5 shifts a week), and to support this as a critical resource would require more staff and/or more automatic tools and ability to trouble shoot most problems from home for off-hours support.

Increased robustness in the tape library

The tape library is neither fault tolerant nor redundant. It goes down when the gripper fails (a few times a year), and our support is normally next business day (3 days if failure occurs on a Friday).

Risk mitigation is to hold all raw data in Lustre until the library is back online. With 3 PB usable and a DAQ rate of less than 1 PB per week, at worst some data that is already in the tape library could be flushed. No loss of data, but could reduce farm utilization.

5. Upgrade to IBM 4500 Tape Library w/ two robots & grippers

Or buy a small 4500 and keep the 3500 running. The 4500 can initialize LTO-7 media to become LTO-M8, saving us $500 / PB. Beneficial, but not enough to cover this cost. We have requested a budgetary quote for the upgrade for planning purposes.

Avoiding "one deep"

With a small staff size for lean operations & support and architecture & procurement it is always a challenge to avoid being one deep. This includes many areas where experience is highly valuable:

- Operations support
  - File systems (especially Lustre)
  - Batch system (slurm)
  - GPU hardware and environment
  - HPC networks (Infiniband, OmniPath)
- System software analysis, design and support
- System architecture
  - Design optimization
  - Procurement (most science per dollar)

We are aware of this challenge, with active efforts to not be one deep in most areas. (Having 2 system architects is too expensive; address via quick replacement when needed as it is rarely more urgent than 4 months.)

Summary

Scientific Computing is doing a great job with a lean staff and a carefully optimized set of resources.

There are many areas where support of scientific computing could be grown and improved. Most would be of a sufficiently high expense that higher level optimizations would have to take into account anticipated costs and benefits to assist in establishing the correct lab priorities.