Data Analysis Coordination, Planning and Funding

Graham Heyes
Data Acquisition and Analysis group.
Physics division, JLab

May 2010

Overview

- Funding
- Organization
- Project management
- Requirements and cost
- Current plans
- Future vision
- Concluding remarks

Projects and funding sources

- 12 GeV project.
  - Funding for hall, detector and accelerator construction.
  - Little or no funding for offline computing resources.
- Nuclear physics general operation.
  - Funding for the operation of the 6 GeV program.
  - Funding for the operation of the 12 GeV program after the upgrade (mind the gap!).
  - Contains most if not all of the planned funding for nuclear physics offline.
- Baseline Improvement Activities (BIA).
  - The lack of funding for computing infrastructure in the 12 GeV project was supposed to be offset by a growth in the 6 GeV budget that never materialized.
  - BIA gets us closer to where we should have been at the end of the 6 GeV program by filling in critical gaps.
  - Computing was dropped from BIA funding.

Organization

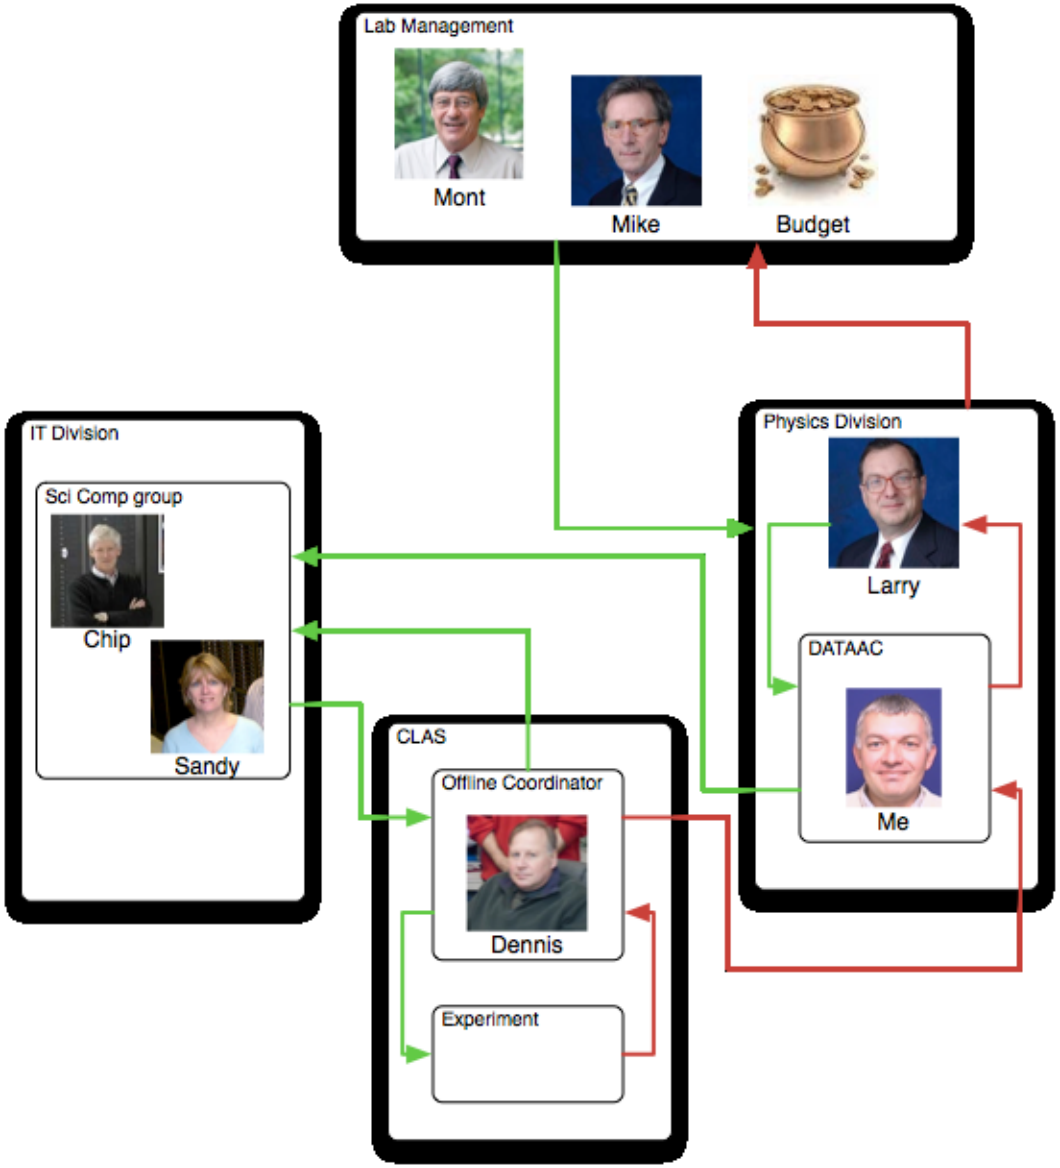

- Each experiment communicates with the hall offline coordinator who grants resources via requests to the SciComp group. Usually via Sandy Philpott who manages day-to-day Farm operations.
- If the coordinator needs a resource that is not already allocated to the hall then he asks me in my role as cross-hall coordinator and I grant via a request to SciComp.
- If I need additional resources not already allocated to Physics I ask Chip what he needs to be able to provide them. If I can't meet those requirements I make a request to Larry (Physics division), usually for funding or staff.
- If Physics division requires additional resources to meet the requirement that is communicated to lab management.
- Why does this matter?
  - You need to talk to the right person to get things done!
  - Instant gratification is not always easy.

Project management

- Project management at JLab
  - Organizations - groups of people and resources able to perform tasks.
  - Projects - groups of tasks.
  - Work plans -annual plans that map projects and tasks onto people and resources.
  - Service provider requests - if my project needs a resource that I don't have I can request it of someone else who tells me what is reasonable and how much it will cost.
  - The work plans are used by lab management to prioritize projects and plan budgets.
- Why are you hearing this?
  - Work plans are written several months before the start of the fiscal year . If you will need something I need to know several months before the start of the fiscal year that you will need it in. (A year or two before is better!)
  - In the current climate walking up to Larry, Mike or Mont and asking for additional funds part way through the fiscal year is certain to fail.

Requirements and estimate inputs

- There are several types of computing:
  - Simulation, Acquisition, Calibration, Reconstruction and Analysis.
- There are four main types of requirement:
  - Long term storage (tape).
  - Compute power to process data.
  - Network bandwidth to move data.
  - Short term storage (disk).
- Maintenance:
  - We need to replace aging hardware at some rate.
- What is the cost of hardware and media.
  - Base on industry trends.

CPU requirements

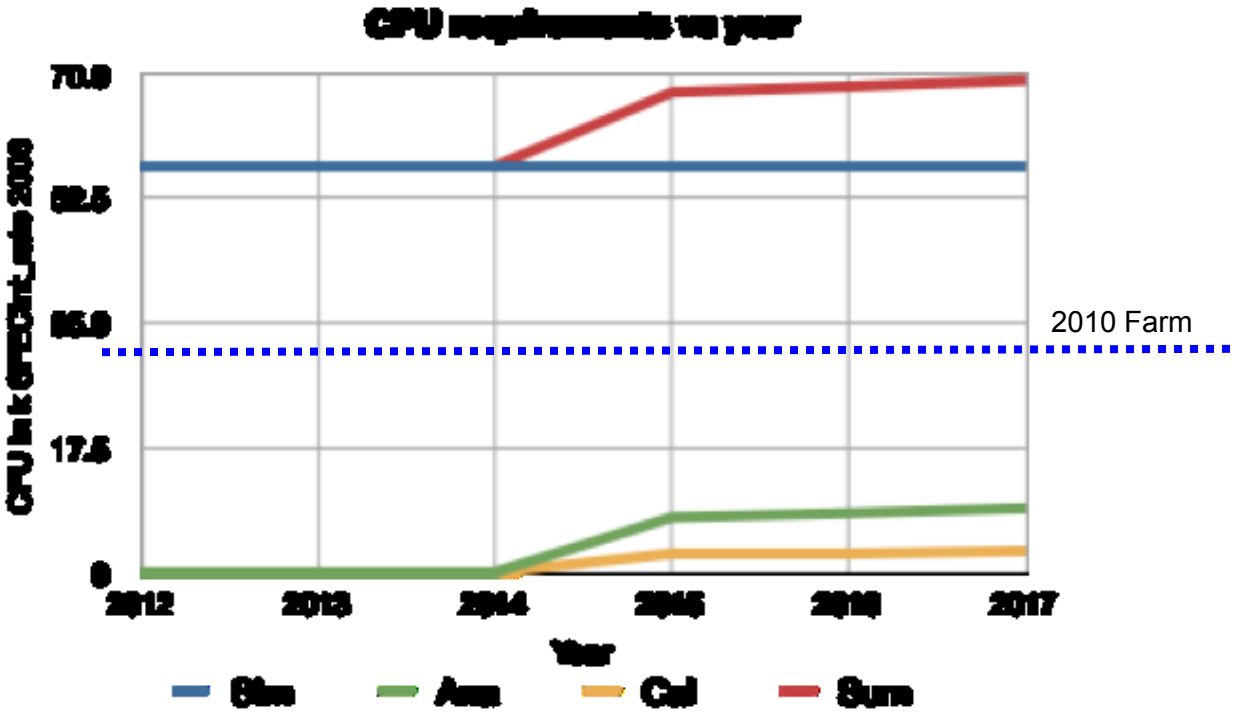

2010 Farm

This chart is based on numbers provided by CLAS12.

The blue line is the processing required to simulate 10^{12} events per year.

Is this a trustworthy number?

How long does a simulation Job take?

How many events are simulated per job?

Performance trends

- See Chip's talk from earlier:
  - The performance of computer hardware has historically increased at a rate that approximates to Moore's law one expression of which is:
  - " The performance of computers per unit cost—or more colloquially, "bang per buck"—doubles every 24 months."

- Measurement of the performance of the last few generations of farm nodes confirm this trend and there is no reason to believe that the trend will break in the next few years.

Requirements into resources

- Summing CLAS12 and GLUEX requirements we need a farm of ~100 x 2009 model nodes in 2012. That is twice the current 64-bit farm. We are on-track to reach that with the current budget of $50k for 24 more nodes each year.

- By 2015 GLUEX pushes the requirements to a farm of 500 x 2009 model nodes but we anticipate 2015 nodes being a factor of 3 to 4 higher in performance. So a ~150 node farm is required in 2015.

- In reality this would probably take the form of the addition of 100 x 2015 model nodes to the 2014 farm.

Current plan - Nodes

- The CPU requirements until 2015 can be met by the current budget of $50k per year for nodes. A bump in the budget to $200k is needed in 2015 to accommodate GLUEX calibration and analysis.
- There is a dip in 2013 because the 100 node farm will be complete but the requirement won't go up for another two years. We can use this to oversize the FY 13/14 farm and lower the FY16 bump.

Current plan disk

- If we assume a disk requirement of 10% of the data taken per year and falling disk prices the 1.6 petabytes of disk bought in 2017 cost about the same as 200 terabytes bought in 2012. The average $55k is pretty close to the current $50k budget.

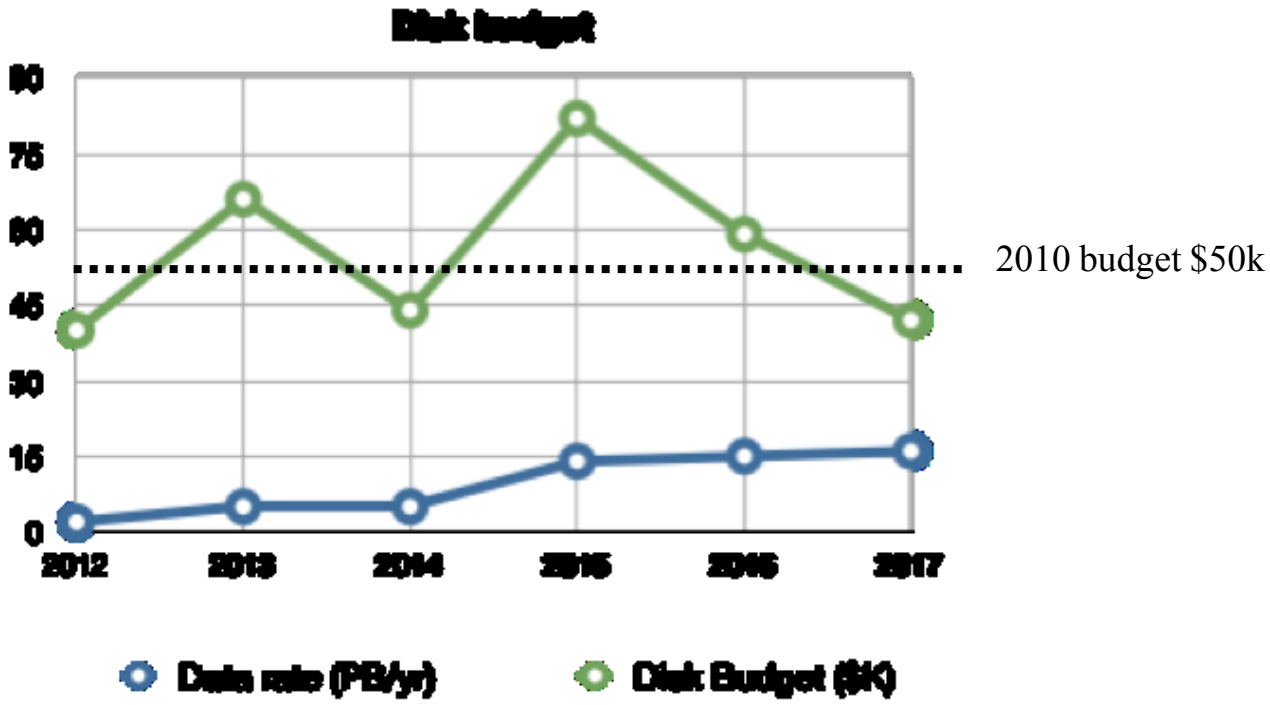

Long term storage requirement

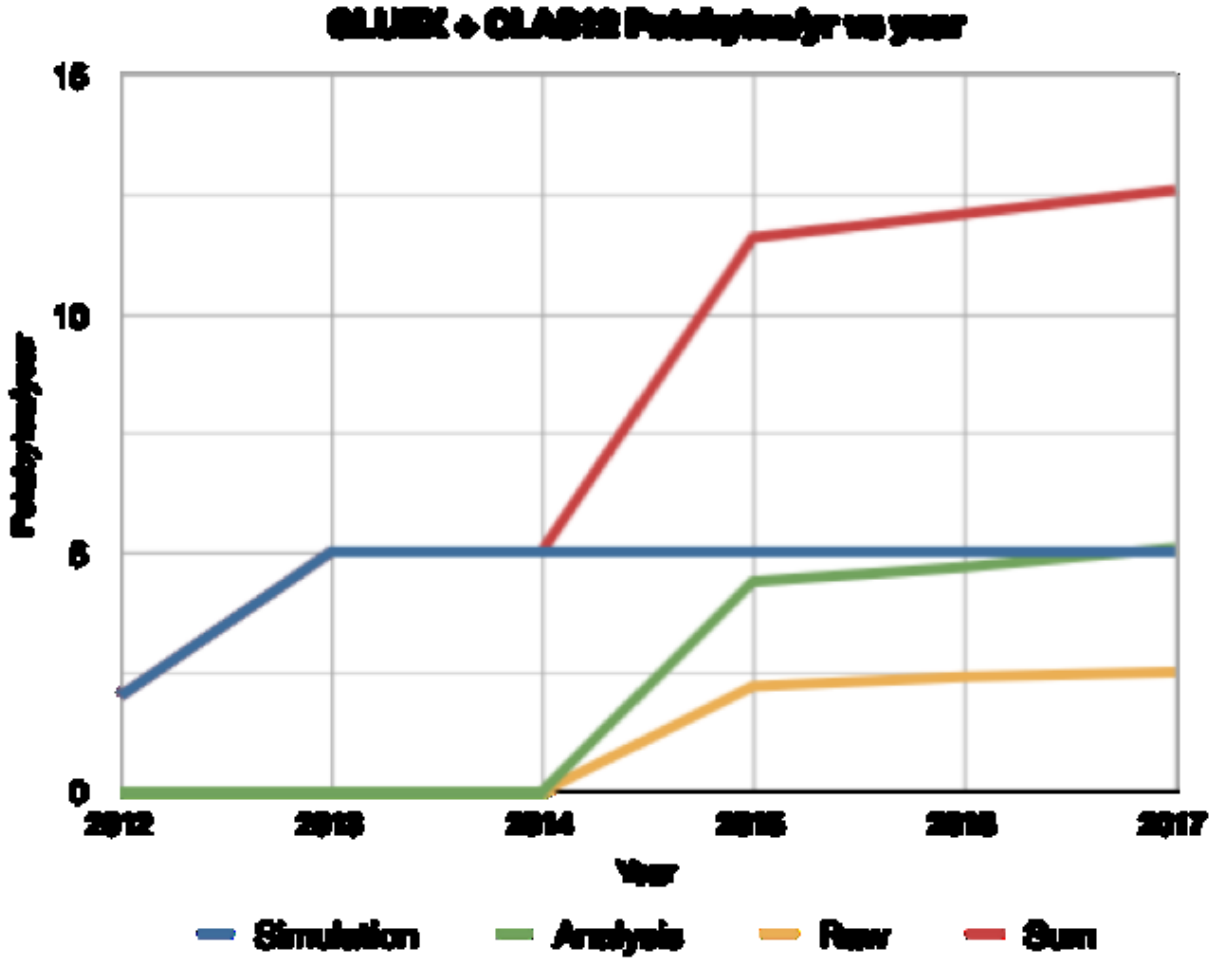

Long term storage cost

- We pay for a library robot, tapes and shelves to put the tapes on.
- The current model is:
  - Media cost is paid by the hall.
  - Tape library hardware and shelf cost is shared by IT and Physics at JLab.
- Although media costs are predicted to fall data rates will rise more quickly.
- Using the current estimates CLAS and GLUEX will generate about 15 PB/Yr -n 2015.
  - In 4 years a library slot may cost $30 and the 4TB tape in that slot $50, i.e. $20 pre TB.
- The long term annual storage budget is therefore ~$0.5M, dominated by CLAS12 - (12 TB/Yr vs 3).
- The cost can be reduced timely reconstruction and removal of data from the library. A third of the $20 per TB is the cost of the library slot and the number of slots per robot is finite.
- Even with tape removal the lab must budget an extra $100k for a second library in 2014.

Future vision

- Annually updated computing plan.
  - An initial computing plan for covering the period from the present to FY17 was completed late in 2009. After much feedback this plan is being updated now.
- Improved coordination and requirements gathering
  - Review and tracing of experiment proposals.
    - The requirements data coming from proposals is currently very poor.
  - Annual update of a short list of experiments expected to require major resources.
  - Cross hall computing coordinator (me) to arbitrate resource requests, coordinate requirements gathering and interaction with IT division.

vision continued.

- Benchmarks
  - Develop “real world” benchmarks based on physics data and code to refine planning.
- Improved communication
  - Meetings and workshops.
    - This meeting is one example, more are planned.
  - Each hall has an offline coordinator.
    - Regular meeting and communication.
  - Set up a website data.jlab.org (still new so not yet ready for prime time).
    - Repository computing plans and requirements.
    - Shared offline documentation
    - Share cross-hall knowledge, tools etc.

Concluding remarks

- Given the current requirements from CLAS12 and GLUEX and current trends in cost and performance we appear to be in good shape.
  - Numbers should be checked! Factors of 2 either way are not good.
- A budget of $50k for nodes, $50k for disk and $50k for other items gets us very close to where we need to be even as late as 2014.
- An additional $200k is needed in FY15 to purchase the additional nodes and disk needed to transition to processing raw data and detector calibration.
- The tape costs level out at about $500k per year in 2016/17 but should start to fall after as the $/TB tape cost improves.
  - Media costs are traditionally paid out of the hall operation budget.
- At 15 petabytes per year the tape library fills up very quickly.
  - The current library outfitted with next generation tape drives would only hold 20 Petabytes!
  - The funding plan must include a $100k library expansion in 2014.
- Even so we must have a good plan for quickly processing the raw data and ejecting tapes from the library to make space for new data.

Last remarks

- There are two main keys to success
  - An understanding of real world requirements and performance.
    - Resource allocation.
    - Planning of future resources.
    - Planning of budget requests.
  - Robust and ongoing communication between groups.
    - Scheduling.
    - Reduce of duplication of effort.
    - Catch challenges before they become problems.

- “The future is not Google-able.”
  - *William Gibson*

- Thank You.