Scientific Computing Allocations Process

August 24, 2016

Overview

The gathering of future computing requirements through external computing reviews and the PAC process will be replaced by a light weight rolling allocations process that provides insights for Scientific Computing into peaks and valleys in the coming 1-2 years' computing requirements of each experiment or hall, and provides users and projects assurances as to the level of resources that will be available for their computing needs.

Introduction and Background

In the past, the Scientific Computing group has tried to gain insight into future computing needs through the experiment proposal process. This has been difficult for users since an experiment might not run for more than 5 years from the writing of the proposal. More recently, 12 GeV Software and Computing Reviews have served as an opportunity every 18 months or so to update computing requirements timelines. Updates to the requirements spreadsheets have also been done driven by specific procurements. We would now like to improve our routine near term allocation and budget planning process in a way that decouples from the reviews and becomes a normal Scientific Computing process, and remove computing from the PAC process.

12 GeV Computing and Storage Growth

In FY2016 and FY2017, the Experimental Nuclear Physics (ENP) computing resources will more than triple compared to FY2015 levels. In August of this year, an additional 44 nodes, each with dual 18 core E5-2697 v4 2.3 GHz "Broadwell" CPUs, are being added to the "farm". Taking into account retiring a handful of 2010 and 2011 nodes, this roughly 50% increase will bring the farm to around 210 nodes with around 4,500 real cores, with up to 7,900 job slots for serial jobs, roughly 50% more than was available before this upgrade.

An additional upgrade of similar size could be done in the Fall to provision for the Spring running period, and perhaps again next summer to provision for the Fall running period. This process of annual doubling might continue into FY2018 and FY2019 as CLAS-12 and other additional high data rate detectors come online.

Storage will also continue to grow at a healthy clip. Usable disk space for ENP grew in August from 600 TB to 800 TB, and will likely reach 1 PB or more by the end of calendar 2016. A growth rate of 40% per year is feasible with constant budgets. The tape library is easily expanded by an additional 20 PB with current LTO-6 tapes, and the cost crossover to LTO-7 media (which is already available) should happen in 2017, allowing even greater growth in the current library. Tape library bandwidth is independently scalable.

High Throughput Farm Computing

JLab's batch computing design and operating model is geared to squeezing as much computing throughput out of limited resources as possible. The hardware is carefully optimized to trade off CPU performance, memory size, disk capacity and bandwidth, and tape library bandwidth (and of course overall system cost) to achieve on average a well balanced system. Projects and users get a "fair share" percent allocation of the resources so that everyone can get work done in each scheduling interval (currently set to a few days). The queuing system allows users to submit more work than can be done in a day or a week, and the system will simply keep processing until all the work is complete. Keeping the queues full enough to occupy all of the computing resources avoids wasting resources.

This of course means that the batch system is sized smaller than the peak load (number of queued jobs at any one time), with submission peaks being smeared out in time to fill in later periods when fewer new jobs are being submitted. Adding more resources (computers, disk, etc.) would make it more responsive (shorter turn around under load), but would lead to there being more idle time on the system during lighter submission rate periods, and would lower the overall throughput per dollar metric of the system (while increasing peak responsiveness).

This time averaging of load is fine over a few days, and for some larger workflows it is fine over a period of weeks or a month. If it takes many months to push through a large load such as a reconstruction run or a simulation campaign, that is not so satisfactory. Dealing with large peaks and valleys that persist for weeks or months thus either requires provisioning above the yearly average load, or requires load shedding or load balancing outside of the farm.

For JLab, LQCD provides a useful load balancing resource. This was especially true when LQCD was several times as large as the farm, and so could serve as a very large load buffer. With this summer's increases in the size of the farm, and LQCD's decommissioning of 2009 and 2010 cluster, resource sharing of conventional nodes shrinks to something closer to 25% of farm capacity in 2016. LQCD is now deploying mostly advanced architectures (Xeon Phi / Knights Landing – KNL – in August of 2016), and it is not yet clear how large a "flywheel" for load balancing LQCD will remain going forward. Testing to determine if Experimental Physics can use KNL will start soon, and could lead to LQCD once again being a 3x sized load sharing buffer.

Without that larger local buffer, there remains the question of handling loads at >125% of farm provisioning for periods of time lasting more than a month.

Load shedding to the Open Science Grid (OSG) is one approach already being pursued by the GlueX collaboration for their simulation campaigns. Using the OSG (or any other grid resource) requires significant investments in knowledge and software development, and relatively smaller investments in hardware. GlueX has already made the learning and software development investments. For JLab to support this for production science, it is still necessary to configure an appropriate amount of disk resources and tape library bandwidth, and necessary to acquire grid computing knowledge in the Scientific Computing group, and this is planned to start in FY2017.

In the future, Cloud Computing resources could also be integrated as a more cost effective solution than local overprovisioning, although in 2016 it is not cost effective unless the cloud is only used for peaks that exist (after LQCD load leveling) for less than 25% of the year. An initial exploration of this space is also planned for FY2017.

New Allocations Process

Because the size of the peaks and valleys drives the future complexity of the computing solution, it is useful to better capture these peaks in as simple a fashion as possible. Further, since requirements evolve as new detectors are commissioned and become better understood, and as accelerator schedules adjust to other pressures, it is useful to routinely update computing requirements – at least once but perhaps twice a year.

It would be difficult to capture 100% of the requirements, as this would involve processing a large number of small requests. Instead, we will assume that the small requests in total evolve only slowly, and thus their load is nearly static in time. Initially we will assume that these small requests will average 20% of our total requirements, and thus we only need to focus upon the 80% coming from larger requests. Expressing requests as core hours (currently normalized to the performance of a 2014 farm node with Haswell processors), the farm at the end of FY2016 will be able to deliver about 40 M core hours per year, and so we are looking to capture any requirement exceeding 1 M core hours.

Simple Requirements

Simple requirements are those with only minor fluctuations month to month. For these, it is only necessary to specify the total requirements for the coming year in units of million core hours. Along with that we would also like to capture something about the utilization of other resources, and so a simple form could collect the following data.

1. number of core hours per year for the coming 12 months (i.e. annual rate)
2. % which is serial, % which is multi-threaded
3. memory footprint of serial jobs and of multi-threaded jobs as appropriate
4. disk footprint needed in terabytes (average for the year)
5. planned number of terabytes to be saved to tape per year

Varying Requirements

If requirements are not roughly constant or slowing growing, then we would like more detailed information at the granularity that matters to the project. For example, if you want to do a simulation campaign in a one month period, you will need to specify requirements at the month granularity. If a one month campaign is broken out as a separate request, then it could be stated as 8 M core hours in one month, to happen (for example) 1 month after the data taking run.

Allocations

One new feature we anticipate is that allocations will include provisioning for known bumps in demand. For example, there might be a commitment to provide N core hours during a Spring running period for near-line data analysis at high priority. Commitments like this might be "firm" for 6 months, and "planned" for later periods, not unlike the accelerator running schedule.

Allocations will be decided by a committee, which may mean that large requests will need some supporting documentation (measured performance of software, basis of estimate such as assumed beam-on time and trigger rates).

This committee might also assist in shaping budgets for the future growth in hardware for local use, and in development of staff and system software to enable load sharing and load shedding.

Transition to the New Process

Scientific Computing is proposing that a small group of people meet as needed over the coming month to define the process sketched above. This group will be tasked with providing recommendations on the makeup of the allocations committee, frequency of allocations updates (e.g. 6 months, 12 months), and thresholds for submission and for more detailed review of large requests. These internal reviews will then replace the external reviews that have been a part of the 12 GeV Software and Computing Reviews.

The committee can also explore whether this process can be used to provide insights into the requirements of the coming 6 months (a first "firm" period).

We will describe the resulting process to what might be our last 12 GeV Software & Computing review committee this coming November, and solicit their comments. Getting the process up and running by then would thus be valuable.

You are encouraged to think of this as a part of the peer review process whereby we all work together to pursue excellence in the Jefferson Lab science program. Yes it will entail some work, but if the process is well defined and made sufficiently light weight, the benefits should far outweigh the costs.