



Procurement Procedures at FNAL

Ian Fisk
LHCC-LCG Review
November 14-15, 2005



Motivation of Procurement Strategy



The goals of the procurement process optimize a number of, sometimes orthogonal, constraints

- ➔ Enable the stakeholder to receive, commission and deploy enough resources in time to meet their goals
- ➔ Make the most efficient use of the financial resources available
 - Resources needs are large and budgets are limited
- ➔ Protect against procuring equipment that either does not satisfy the requirements, cannot be operated with the existing effort constraints, or is otherwise unreliable
 - It is necessary to discover bad equipment before hundreds of them are at the site
- ➔ Procure equipment that can be deployed given the realities of facility infrastructure and the requirements of other stake-holders
 - Power and cooling have become a limiting resource. Most sites support several activities.



Procurement Process (1/3)



At Fermilab the procurement process runs through several steps

- ➔ The first is to select qualified vendors for equipment
 - Periodically vendors are invited to send a reference system that is chosen to resemble the current stakeholder requests to FNAL
 - These systems are evaluated on how well they meet the specs, quality of assembly, temperature issues, etc
 - The last time this was performed approximately 25 vendors participated, not all passed. Eliminates vendors who cannot support
- ➔ The vendors that are selected from the qualification processes are invited to submit a price performance bid.
 - There are 5-6 vendors left to compete on FNAL contracts
 - FNAL is in the process of refining the procurement process and there are discussions of decoupling qualification and price performance
 - The goal is to keep the number of vendors responding to procurement requests small enough to perform detailed evaluations without infinite people
 - While keeping the responding number large enough to get competitive bids



Procurement Process (2/3)



At the time of a procurement, a stake-holder will put in a request to the lab. Systems can be received in a whole order or in quanta of racks

- ➔ Vendors are given two weeks to respond with a bid for the full request and a single evaluation unit
- ➔ The evaluation units are tested by a team led by the stake-holder for a period of one week
 - The system is compared to the bid requirements and the enclosure to facility constraints
 - A benchmarking suite to determine performance
 - Use the standard batch environment to replicate running conditions
 - The power under load conditions is measured
 - The stakeholder returns the list of technically acceptable bids to the procurement office
 - Only the units with a chance of winning are carefully evaluated
- ➔ The chosen vendor is given 4 weeks to deliver the systems to Fermilab



Procurement Process (3/3)



After assembly on the FNAL site the systems enter two weeks of burn-in

- ➔ The burn-in procedure is defined in the bid request. This currently consists of
 - SETI at home for the CPU, Bonnie++, and a memory test application
 - These typically discover significant system failures.
 - Separating benchmarking and burn-in: benchmarks are performed using stakeholder applications while burn-in is done using common applications
 - For an order to be accepted the uptime on the cluster must exceed 98% uptime for the two week burn-in

After burn-in the systems are available for production use. The vendor invoices are paid.



Recent US-CMS Experience



US-CMS has moderately large procurements to make in preparation for the start of the experiment.

- ➔ In order to avoid the liability associated with procuring all the nodes simultaneously, which requires a lot of effort to commission and risks having an entire year's procurement be unacceptable, CMS experimented with taking possession in stages.
- The vendor is notified of the US-CMS desire for equipment purchases for the year
- The procurements are set up as options
 - US-CMS is committed to buying a single rack. When the rack has completed the burn-in process the option to procure additional racks can be exercised
 - In the 2005 procurement a buy of 280 nodes was performed in 3 options (1 racks and 2 orders of 3 racks)
- The process is extended, but the liability is limited
- The effort is spread over longer, but we can focus on incoming systems
- Surprisingly the vendors seemed to like the process. Receiving also
- System works because the vendor trusts that we will buy the units



Timeline



At the moment the time from order to gear on the floor is 10-11 weeks assuming no problems

- ➔ 1-2 weeks to prepare bids
- ➔ 2 weeks to receive responses and evaluation nodes
- ➔ 1 week to evaluate nodes
- ➔ 4 weeks to receive systems
 - Hardware addresses and local IPs are received in advance
- ➔ 2 weeks of burn-in

FNAL has been working to reduce the time lag

- ➔ Initial burn-in period was 30 days, which was determined to be unnecessarily long
- ➔ Hard to squeeze current periods and maintain the existing structure

The universities at the Tier-2 centers are generally operating on shorter time scales, but generally do so by eliminating steps that increase the risk of a bad order.



CPU Ramp to 2008

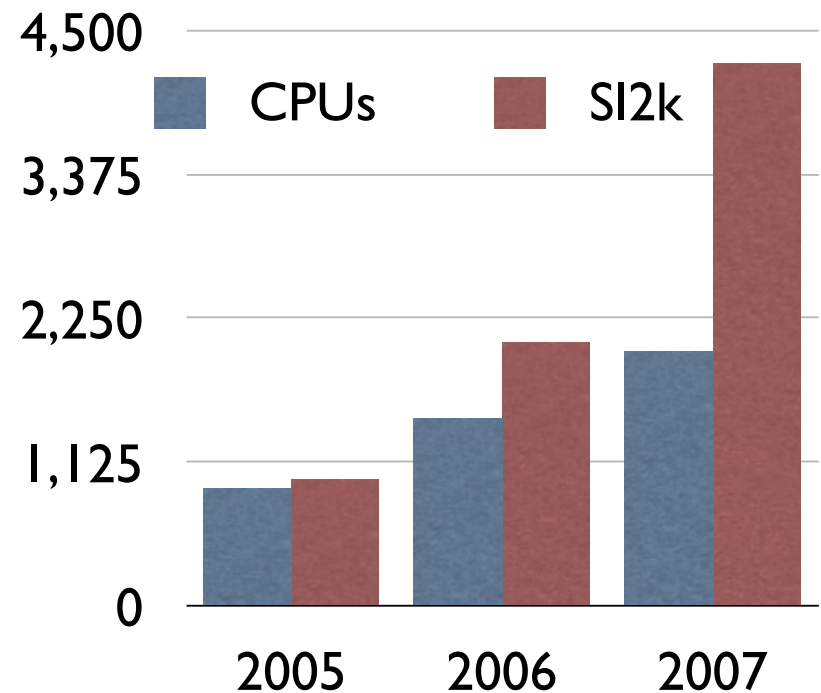


The operational ramp to the start of the experiment is manageable for processing resources

- ➔ Experience at FNAL configuring and running farms this size already for Run2

The increase in number of nodes is almost linear

- ➔ Performance increase is a fairly conservative improvement estimate
- ➔ Dual cores CPUs may improve the situation
- ➔ Ramp of power and cooling are a significant issue





Disk Ramp to 2008



While FNAL has several hundred TB under management for the Run2 experiments, CMS dCache space at FNAL is ~100TB in 2005

Very steep operations ramp in disk storage before the experiment start

- ➔ Approximately a factor of 20
- ➔ Expecting to gain from improvements in capacity, benefit from later acquisitions. Weighed against need for operations experience.
- ➔ We are scaling dCache well past our current operational experience.

