

We will be having our regular bi-weekly Infrastructure WG telecon on Monday, December 21, at 12 Noon CT.

Agenda

- HPC Allocation Update
 - ◆ Current Allocation (as of Dec20)
 - ◇ Allocated: 30K SUs on abe; 30K SUs on lincoln (nVidia Tesla GPUs)
 - ◇ Used: 107 SUs (abe); 0 SUs (lincoln)
 - ◇ Remaining: ~29K SUs (abe); 30K SUs (lincoln)
 - ◇ How to get access: contact MikeF
 - ◆ Next TG Allocation Cycle
 - ◇ Proposals due Jan15 for Apr1 allocations
 - ◆ Note the following:

The TeraGrid will host a teleconference on two different dates to provide step-by-step instructions and answers to your questions on how to write and submit a successful proposal.

* **Allocation Request Guidelines Telecon 1: Dec. 22, 1:00-2:30 CST (central)**
* **Allocation Request Guidelines Telecon 2: Jan. 4, 1:00-2:30 CST (central)**

Call 1-866-740-1260 and enter Access Code 8229741.

To view the presentation online during the talk, visit "readytalk":<http://www.readytalk.com/> during the call and enter the same access code in the _Participant Login box_. (For our records please enter your name, email and institution when requested.)

- Cost Sheet Update
 - ◆ Baseline version is v45
 - ◇ [?https://www.lsstcorp.org/docushare/dsweb/Get/Version-12185/Infrastructure-Costs-v45.xls](https://www.lsstcorp.org/docushare/dsweb/Get/Version-12185/Infrastructure-Costs-v45.xls)
 - ◇ caveats apply: v45 does not *exactly* match PMCS
 - ◆ Current version now v69
 - ◇ [?https://www.lsstcorp.org/docushare/dsweb/Get/Document-6284/Infrastructure-Costs-v69.xls](https://www.lsstcorp.org/docushare/dsweb/Get/Document-6284/Infrastructure-Costs-v69.xls)
 - ◇ [?https://www.lsstcorp.org/docushare/dsweb/Get/Document-8189/CostSummaryWithBaseline-](https://www.lsstcorp.org/docushare/dsweb/Get/Document-8189/CostSummaryWithBaseline-)
 - ◆ Summary of Changes
 - ◆ Questions & Notes
 - ◇ Ramp up: One of the things in the cost sheet that I wonder about is our "ramp up", i.e. we're currently planning on buying 1/3 of the hardware 3 years early, 2/3 two years early, etc. I wonder if 3 years early is a little too soon.
 - ◇ [?LSST-9](#) Estimating the cost of forced sources
 - What about increased CPU requirements to process forcedsources (above and beyond what would be required without forcedsources)?
 - ◇ [?LSST-65](#) UPS
 - No UPS in the PCF; What are LSST requirements for UPS? Do we need one? If so, how much time do we need?
 - Possible construction implications.
 - 85 degrees
 - ◆ Upcoming Changes
 - ◇ **Priority is getting a new BaseFloorSpace.xls to RonL/JeffB**, which depends on:
 - [?LSST-50](#) Floorspace tab: Floorspace calculation does not take into account the increase in drive capacities over time
 - [?LSST-40](#) Disk Cost/Capacity? Trends (3 yr step, etc.)

- ?LSST-71 Compute model currently based on Rpeak. This needs to be changed. Rmax is better, but still not right. What to use?
- ?LSST-55 How are TFlops mapped to nodes? How does this evolve over time? [constant TF/core; constant \$/node; moore's law for cores/node (model now more flexible)] (\$-451K/\$-4096K)
 - Radical reduction in number of compute nodes. Arch from 741->2049 to 99->102; Base from 337->376 to 48->49
 - ?LSST-80 Planning to introduce a "Factor X" for the out years
- ◇ ?LSST-72 Update PMCS Baseline in Cost Sheet
- ◇ ?LSST-78 Move the 3% CPU spare from document 2116 "CPU Sizing" to document 6284 "Cost Estimate"
- ◇ ?LSST-79 Add tape library replacement to ArchAOS and BaseAOS
- ◇ ?LSST-10 Update Power & Cooling at Base Site (info already received from RonL)
- ◇ ?LSST-47 Power Costs at BaseSite: Use Historical Data to Model Future Power Prices
- ◇ ?LSST-36 Update Power & Cooling at ArchSite
- ◇ ?LSST-28 Optimal CPU Replacement Policy
- ◇ ?LSST-36 P&C and Floorspace at PCF (rates, payment approach, green features of PCF)
- ◇ ?LSST-69 New Model for Floorspace (Lease Costs) at ArchSite
- ◇ ?LSST-14 Processor Sizing Update (Doc2116 LSST CPU Sizing)
- ◇ ?LSST-37 Missing controller costs for disk
- ◆ Next steps with cost sheet
 - ◇ Full review each of the elements of the cost sheet (boxes of the mapping document)
 - More readable description of the formulas being used
 - Identification and documentation of assumptions
 - Identification and documentation of external data input
 - ◇ Serves two significant purposes
 - Allows for better internal reviews (validation of models and information used)
 - Provides justifications for external reviews
 - ◇ Results in an updated (or replacement of) Document-1684 and related documents ("Explanation of Cost Estimates")
- DC3b Storage Options/Costs?
 - ◆ [?http://dev.lsstcorp.org/trac/wiki/DC3bHardwareRequirements](http://dev.lsstcorp.org/trac/wiki/DC3bHardwareRequirements)
 - ◆ ?LSST-11 DC3b Hardware
 - ◆ Request in to TG allocations of PT1
 - ◆ Request in to NCSA allocations for PT2, PT3
 - ◆ Discussion with Allocations
 - ◇ commitment of **10TB spinning disk from TG** (covers PT1)
 - ◇ tape complicated - can't commit until ~Jan
 - old tape system/new tape system; tape format compatibilities; funding sources & HPC followons; etc.
 - ◆ Spinning Disk
 - ◇ Should be doable if we stay close to the lower end of the range
 - ◆ DB Storage
 - ◇ Can use existing SAN
 - ◇ ?LSST-73 Add additional 3.3TB of our existing SAN allocation to lsst10 /scr
 - ◆ Tape
 - ◇ This is going to be difficult

- ◇ \$62/TB (for single copy) [\$25/tape=400GB]; for 300TB is \$19K
- ◇ Q1: *Can we delete PT1 data at the start of PT2, and PT2 data at the start of PT3?*
How much tape space would such a policy save us?
 - Given the anticipated costs for tape storage, this could save quite a bit of money.
 - But, this reverses previous discussions/requirements about serving PT1 data during PT2, and PT2 data during PT3.
- ◇ Q2: *Is data loss acceptable?* (750 tapes)
 - I have a question in right now regarding tape failure rates
- ◆ Compute
 - ◇ SU being addressed by our first topic above
- DC3b Data Serving Options/Costs?
 - ◆ [?http://dev.lsstcorp.org/trac/wiki/DC3bDataServingRequirements](http://dev.lsstcorp.org/trac/wiki/DC3bDataServingRequirements)
 - ◆ [?LSST-54](#) Connections Speeds between lsst10 and the SAN Storage. We need 300 MB/s. What are our options?
 - ◆ Image Retrieval needs are unspecified (servers? spinning disk?)
- Distributed File Management (REDDNET/Lstore/iRods/DataNet) iff we have the right people on the call
- 11th LCI International Conference on High-Performance Clustered Computing
 - ◆ March 8-11, 2010 Pittsburgh
 - ◆ [?http://www.linuxclustersinstitute.org/conferences/](http://www.linuxclustersinstitute.org/conferences/)
- InfraWG Ticket Update

Notes

Attendees: K-T, Arun, Ray, MikeF

- Action items reflected in JIRA tickets.
- UPS is more important at base site, i.e. the UPS requirements likely won't be the same at the two sites
- DC3b PT data is cumulative, i.e. cannot dump PT1 data during PT2, etc., to save space/money
- A new diskIO sheet is coming, with a summary tab formatted to match previous versions
- Get Arun a JIRA account

Useful Links

- InfraWG Home Page
 - ◆ [?http://dev.lsstcorp.org/trac/wiki/InfrastructureWG](http://dev.lsstcorp.org/trac/wiki/InfrastructureWG)
 - ◆ [?All InfraWG Tickets](#) (in priority order)