We will be having our regular bi-weekly Infrastructure WG telecon at a special time, Wednesday, February 17, at 2P CT (12N PT).

Agenda

- Access to mss for ImSim team

- Existing Resource Usage Update (as of Feb16)
    - TeraGrid Resources (Startup Allocation)
        - Service Units
            - Abe: Allocated: 30K SUs; Remaining ~29.4K SUs
            - Lincoln (nVidia Tesla GPUs): Allocated: 30K SUs; Remaining 30K SUs
        - Disk Storage
            - Allocated: 5TB; Remaining: 5TB
        - Tape Storage
            - Allocated: 40TB; Remaining: 40TB

- HPC and MSS File I/O Performance Benchmarks
    - reading from abe filesystems
        - abe:/cfs/projects/lsst -> /dev/null = 150MB/s (lustre; dd on worker node)
        - abe:/scratch/batch -> /dev/null = 146MB/s (lustre; dd on worker node)
    - copying on abe
        - /cfs/projects/lsst -> /scratch/batch = 31MB/s (lustre to lustre; dd on worker node)
    - mss to abe
        - mss -> abe:/cfs/projects/lsst = 144MB/s (uberftp active)
        - mss -> abe:/cfs/projects/lsst = 144MB/s (uberftp active parallel 4)
        - mss -> abe:/cfs/projects/lsst = 140MB/s (uberftp active parallel 8)
        - mss -> abe:/cfs/projects/lsst = 143MB/s (uberftp active parallel 8 tcpbuf 4000000)
    - lsst9 -> mss
        - lsst9 -> mss = 34MB/s (scp)
    - mss tape to mss cache
        - 277MB/s ("1TB/hr" is SET quote - current mss system)
    - Bottom line: **30 seconds per image** (assuming ~4GB files, compressed)

- Mass Storage Symposium ?Science Driven Data Management?, May 4 ? 5, Lake Tahoe
    - DM invited to present, need to send title(s)
    - Arun, Jacek, Ray - submit a paper title

- Update on LSST Database Performance Tests Using SSDs (Arun/Jacek?)

- Update on Lawerence Livermore database scaling test

- DC3b Infrastructure for the Performance Tests
    - ?http://dev.lsstcorp.org/trac/wiki/DC3bHardwareRequirements
    - ?LSST-11 DC3b Hardware
    - Compute
        - Research Proposal submitted to TeraGrid for allocations starting Apr1
        - ?https://www.lsstcorp.org/docushare/dsweb/Get/Document-8529/LSST-TeraGrid-Proposal.pd
        - should hear back in early March

- Database Disk
  - ?LSST-73 Add additional storage from our existing SAN allocation to lsst10
    - We are adding our remaining SAN allocation to lsst10. The total space available for MySQL data will be 17.3TB. See ?LSST-73 for more details.
    - This is happening Monday Feb 22 during the outage.
- Tape Storage
  - Total (raw) tape needed is (203TB, 288TB, 449TB)
  - The **tape gap** is 0 for PT1, 88TB for PT2, **249TB** for PT3 (contingent upon 200TB raw from TG)
  - Pricing
    - $62/TB (for single copy) [$25/tape=400GB]; for 300TB is ~$19K (LTO-3)
    - $31/TB (for single copy) [$50/tape=1.6TB]; for 300TB is ~$9K (LTO-5) [plus faster bandwidth than LTO-3]
    - Note1: Working with PI for possible purchase strategies, which include PI subsidizing during tape usage and LSST subsidizing LTO-5 tapes, to avoid the need to buy LTO-3 (old technology) tapes.
  - Additional notes
    - Note2: Mass storage will exist post-2010 (ongoing talks with NCSA PI); New system by Oct1 (estimated)
    - Note3: Estimated DC3b data loss from tape failures expected to be 2-3% per year

- DC3b User Access
  - DC3bUserAccess
  - **Development of Use Cases**
  - ?LSST-54 Connections Speeds between lsst10 and the SAN Storage. We need 300 MB/s. What are our options?
    - **Do we really need 300MB/s**? (Jacek)
    - Adapter slots on lsst10 will not support 8Gb HBA
      - in the process of getting price estimates for a new database server

- REDDnet Update

- Mass Storage Access Requirements
  - Do we need access to mss either to or from any lsst* machine or ds33?
  - lsst10 for catalog backups / replication to REDDnet?

- Database access from Pipeline processing during DC3b
  - discuss details, specifically, is access occuring:
    - during pipeline startup?
    - during stage pre/post process?
    - during stage process()?
  - effects on performance, scalability, etc.

- Cost Sheet Update
  - Baseline verion is v45
    - ?https://www.lsstcorp.org/docushare/dsweb/Get/Version-12185/Infrastructure-Costs-v45.xls
    - caveats apply: v45 does not *exactly* match PMCS
  - Current version now v74
    - ?https://www.lsstcorp.org/docushare/dsweb/Get/Document-6284/Infrastructure-Costs-v74.xls
    - ?https://www.lsstcorp.org/docushare/dsweb/Get/Document-8189/CostSummaryWithBaseline-

- Summary of Changes
  - ?LSST-94 Floorspace tab: Increase rack depth from 3.0 to 3.5 (+50sf both sites) (+19K/+153K)
  - ?LSST-95 Floorspace tab: Add calculation for gross floorspace for the base site
- 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.
- Upcoming Changes
  - **Priority is updating the Power & Cooling estimates**
    - ?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-36 P&C and Floorspace at PCF (rates, payment approach, green features of PCF)
  - ?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-28 Optimal CPU Replacement Policy
  - ?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")

- InfraWG Ticket Update

Notes

Attendees: RobertL, TimA, TomH, AndyC, GarrettJ, JacekB, JohnP, RayP, DeborahL, DickS, SuzieD, SchuylerV, DanielW, ArunJ, JeffK, KTL, MikeF

- Preliminary discussion regarding importing ImSim data into the DC3b storage
  - MikeF will coordinate followup meeting to pursue details
- DB Testing with SSDs - in setup phase - user accts, etc.
- LLNL DB Scalability Testing - working out issues from 25 node test - 100 node test upcoming
- DC3b Infrastructure -- adding more disk storage to lsst10 database server on Monday
- DC3b User Access use case discussion - lots of good discussion - abbreviated notes follow - send errors,omissions,corrections to MikeF

        - scope is DC3b only; limited resources; focusing on simplest set of interfaces that meets requirements
        - initial set of use cases: sql interface to catalog; http interface to image files; web interface to catalog schema information
        - baseline modus operandi: use of scripting with SQL and WGETs; sample scripts can be provided
        - additional functionality
            - image cutout service
            - bulk upload to db
            - developer/facilitate testing
        - web page interface
            - IPAC's plate
            - interface to scripts
            - portals from other projects
            - VO tools
- Action items reflected in JIRA tickets.

Useful Links

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