Computing Update

A. Lazzarini

LSC Computing Committee

LIGO-G030396-00-E

LSC Computing Update
ITR 2003 Proposal

- ***Grant awarded to ITR2003 Proposal***
  - » Proposal for 3 key tasks to complete deployment, support operation of Tier 2 Centers
    - – Provide authenticated and secure access to LIGO data by collaboration members by developing a Tcl API to the Globus Package for use with existing infrastructure
    - – Continued development & deployment of LIGO Data Replicator (LDR)
      - Efficient and secure mirroring of critical datasets across the LIGO Data Grid.
    - – Port DMT software to a grid-based computing model
  - » Proposal requests 6 FTEs for 4 years
  - » 4 year award for $3M
  - » FTE allotment:
    - – 1.5@CIT
    - – 1@MIT
    - – 1.75ea@UWM, PSU

LSC Computing Update
Software Releases

- **S3 Release: LDAS 0.8.0 - Mid October 2005**
  - » Planned version for S3 Run
  - » New threaded diskCacheAPI and metaDataAPI
  - » ~100% performance improvement
  - » Fully ported to Redhat 9
  - » Based on updated suite of packages in the LDCG directory
- **LDAS 0.9.0 - Late January 2004**
  - » Added functionality to address any issues discovered in S3 Run
  - » Continuing performance boost
  - » Better support for foreign Frame structures (non-FrameBuilder Frames)
  - » Greater reliability (lowering of memory leaks and API restart rates)\
  - » Concentrated effort on bug fixes
- **S4 (after seismic upgrade) Release: LDAS 1.0.0 - ~ July 2004**
  - » Raise reliability levels in anticipation of 6 month runs
  - » Continued optimizations through profiling methods
  - » Upgrade to newer packages in LDCG
  - » Upgrade to newer Linux and Solaris OS's if necessary
  - » Conclusion of effort to address identified bugs (hopefully!)

LSC Computing Update

Computing & Archiving - LIGO Laboratory

- **Compute Cycles**
  - » "X" = 70 dual-2.66GHz P4-Xeon nodes with 2GB of RAM, 200GB of disk, and GigE.
  - » LLO = 1X; LHO = 2X; CIT = 3X = 420x2 total CPUs
  - » MIT = 112 single-2GHz P4 nodes with 512MB of RAM and FastEthernet.
  - » ***Aggregate CPU cycles: 1,341 GHz (1.3 THz)***
- **Archive Storage**
  - » LHO, LLO each have 140TB of tape storage that is being brought on-line this month running SAM-QFS. This will hold about 1 year of full frames, and we may "cross-pollinate" the Observatories with each others Level-1 RDS data as well.
  - » CIT has 500TB of tape storage (that can expand to 1PB) also running SAM-QFS.
  - » We are migrating old HPSS data to SAM-QFS for easier access (S2 trend frames done, old E-runs progressing), and re-archiving earlier test SAM-QFS data to the new higher density tape drives.
  - » Disk storage ~ 140TB across all sites by S3

LSC Computing Update

Computing & Archiving - PSU

- 8/02-3/03: trade studies
  - » UMP vs. SMP nodes; interconnects, IDE vs. SCSI vs. FC; h/w vs. s/w NFS
- 3-7/03: negotiation with vendors
  - » Apple, Apro, BlueArc, Dell, HP, IBM, Sun, …
- 7/03: POs placed
- 7-8/03: Equipment arriving, under assembly
  - » “First Cycles”: 9/2/03

- Configuration
  - » Compute nodes
    - – 128x2 2.8 GHz Xeon, 2 GB, 533 MHz FSB
    - – 28x2 3.06 GHz Xeon, 2 GB, 533 MHz FSB
  - » Storage
    - – 34.2 TB striped across 18x14 bay SCSI raid enclosures
    - – 146 GB 10K RPM Ultra 320 SCSI disks
    - – Served by 9x2 3.06 GHz Xeon, 4 GB, 533 MHz FSB
  - » Interconnect
    - – Copper GigE across all; Dolphin (Scali) across 165 nodes

LSC Computing Update
Computing & Archiving - UWM

- 300-node cluster operation
  - » 24/7 with weekly scheduled maintenance for ~ 4 hours
  - » All S1 and S2 RDS data are on disk cache
  - » ~ 30 non-UWM LSC user accounts. (Web page account request.)
  - » Both LDAS and Condor batch systems available in general
  - » User and maintainer documentation located on web page:
    http://www.lsc-group.phys.uwm.edu/medusa
  - » LDAS currently down due to installation problems with 0.7.0

Cluster CPU usage summary

Wed, 13 Aug 2003 11:55 CST

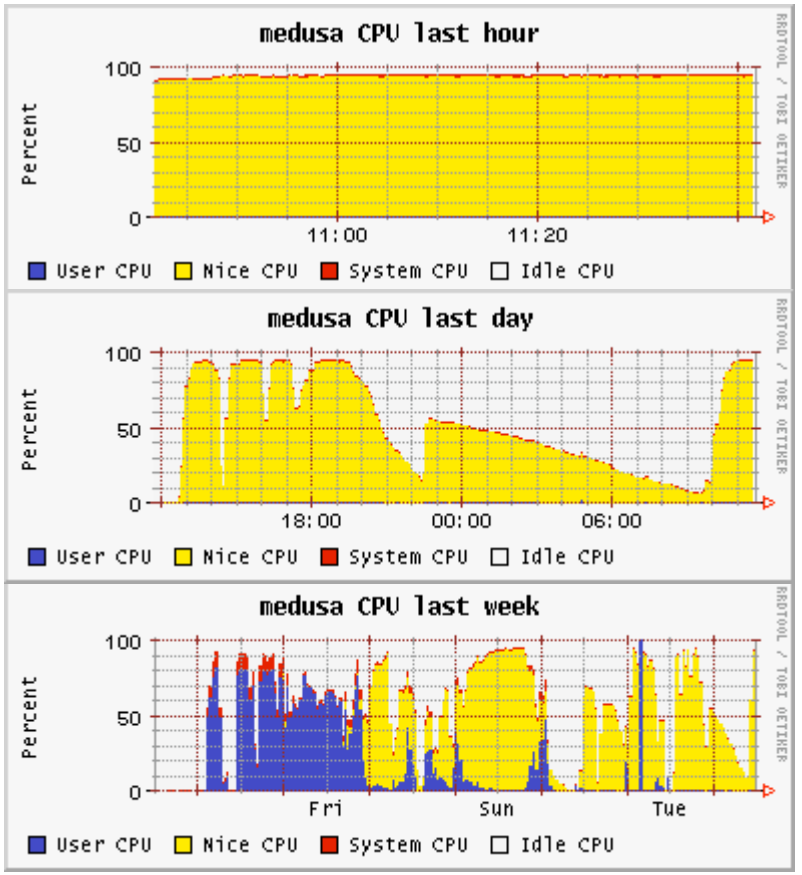

LSC Computing Update
Issues/Plans

- **Preparation for S3: RDS generation**
  - » Developing plan to ensure smooth generation & propagation of RDS dataset(s)
    - – Data set definitions (more than 1 flavor?)
    - – Generation at Observatories, Tier 1, Tier 2 centers
      - Immediate availability at LHO, LLO
      - Maximum delay of 2 weeks for availability at Caltech, then stream to Tier 2 sites - **WAN bandwidth does not yet support streaming site data to Caltech**
    - – Philip Charlton to organize, lead effort, form team including Detector Characterization people, Tier 2 center support
    - – Mock challenge planned for early/mid October to validate plan

LSC Computing Update
Grid Activities (GriPhyN, iVDGL)

- **Preparation for Supercomputing 2003 ("SC2003") Convention**
  - » Scaling up SC2002 demonstration to support large-area unbiased periodic source search
    - – Ultimate goal: use O[10X] more resources than LSC can bring to bear on the problem => 5000 - 10000 CPUs for 30 days (**!!**)
    - – Realistic near term goal: 1000 nodes
  - » Joint effort with CS team at USC/ISI, and LSC CW working group -- Caltech, UWM, AEI
  - » Have asked advice of iVDGL, GriPhyN management for access to O[5k] machines for 1 month
    - – Response not promising
    - – We will scale analysis to as many CPUs as we can scavenge
    - – This effort intentionally designed to *stress* the grid "establishment"

Pull back the curtain to see what the 'Wizard of Grid' looks like …