

Louisiana Tech Site Report

DOSAR Workshop IX
in South Africa

April 8, 2010

Michael S. Bryant
IT Systems Administrator



LOUISIANA TECH
UNIVERSITY®



Louisiana Tech University, LONI, and PetaShare

COMPUTING IN LOUISIANA



Current Tier 3 Resources

- Small Tier 3 workstation (T3w) cluster
 - Two 8-core Xeon 5300/5400 RHEL5 workstations (16 cores, 32 GB/24 GB)
 - Old dual-processor Xeon compute nodes (8 CPUs, 1 GB/core)
 - Repurposed for hosting DOSAR's website/VOMS and local Physics department servers
 - Several Pentium 4 graduate student desktops
- Tier 3 storage
 - Shared work area (NFS): 500 GB
 - Performance issues and crashes when pulling large datasets using dq2-get
 - Each workstation has 250 GB locally, while desktops have ~10-80 GB each
- Software available
 - ROOT, DQ2 client, OSG client, and Athena
- This, of course, does not meet U.S. ATLAS Tier 3g requirements
 - > 80 cores
 - > 20 TB of disk space



Louisiana Optical Network Initiative



LONI is an advanced fiber optic network and state initiative devoted to enhancing research by offering state-of-the-art computing and storage technologies throughout Louisiana.

The **LONI Institute** uses LONI to drive research and education within the state and create a multi-disciplinary collaboration of researchers focused on advancing *computational sciences* in the state.

- A **40 Gb/sec fiber-optic network** connected to the National LambdaRail (10 Gb/sec) and Internet2
- Provides 12 high-performance computing clusters around the state (11 of which are online)
- Over **85 teraflops** of computational capacity
- Around 250 TB of disk storage for research data

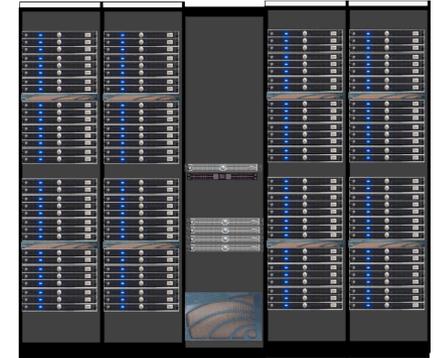


MOST POWERFUL
DISTRIBUTED SUPERCOMPUTER RESOURCE
AVAILABLE TO ANY ACADEMIC COMMUNITY
<http://www.loni.org>



LONI Computing Resources

- 11 operational HPC clusters available:
 - Dell 50 Teraflop Intel Linux cluster (*LSU*)
 - 668 compute nodes (**5,344 cores**), RHEL4
 - Two 2.33 GHz quad-core Intel Xeon 64-bit processors
 - 8 GB RAM per node (1 GB/core)
 - 192 TB Lustre storage
 - Five Dell 5 Teraflop Intel Linux clusters (*LSU, LaTech, ULL, Tulane, UNO*)
 - 128 compute nodes (**512 cores**), RHEL4
 - Two dual-core 2.33 GHz Xeon 64-bit processors
 - 4 GB RAM per node (1 GB/core)
 - 12 TB Lustre storage
 - Five IBM Power5 AIX supercomputers (*LaTech, ULL, Tulane, UNO, SU*)
 - 13 nodes (104 CPUs), AIX 5.3
 - Eight 1.9 GHz IBM Power5 processors
 - 16 GB RAM per node (2 GB/processor)



Dell Linux Cluster

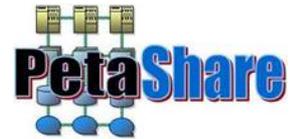


IBM Power5



Statewide Data Grid: PetaShare

- **PetaShare** is a distributed storage management infrastructure that delivers a **global file system namespace to LONI researchers** with advanced and automated data movement capabilities for large-scale data collections.
 - Assembles **300 TB of disk storage** distributed across the state
 - 400 TB of tape storage is also available
 - Leverages LONI's high-speed 40 Gb/s fiber network
 - Based on the Integrated Rule-Oriented Data System (**iRODS**)
 - PetaShare is an NSF funded project led by Tevfik Kosar at LSU
- Initially, **10 TB is devoted** to local U.S. ATLAS analysis
 - Used primarily for long-term storage of analysis data
 - Optimized for use on LONI clusters, but is accessible from other clusters, desktops, and laptops



<http://www.petashare.org>



PetaShare (left) and Painter (right)



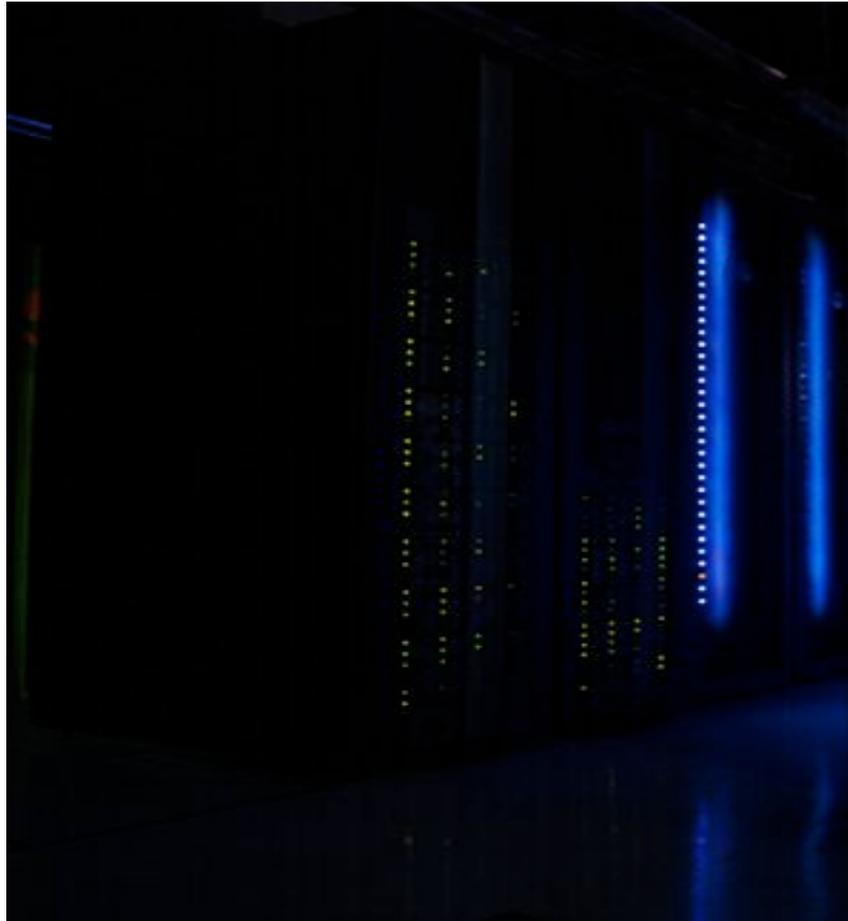
Of the 300 TBs of PetaShare storage, 40 TB sits right here on the Louisiana Tech campus.



Six of these 128 node (512 cores) compute clusters are available to all Louisiana researchers around the state. The LONI cluster [Painter](#) is located at Louisiana Tech University in the Data Replication Center.

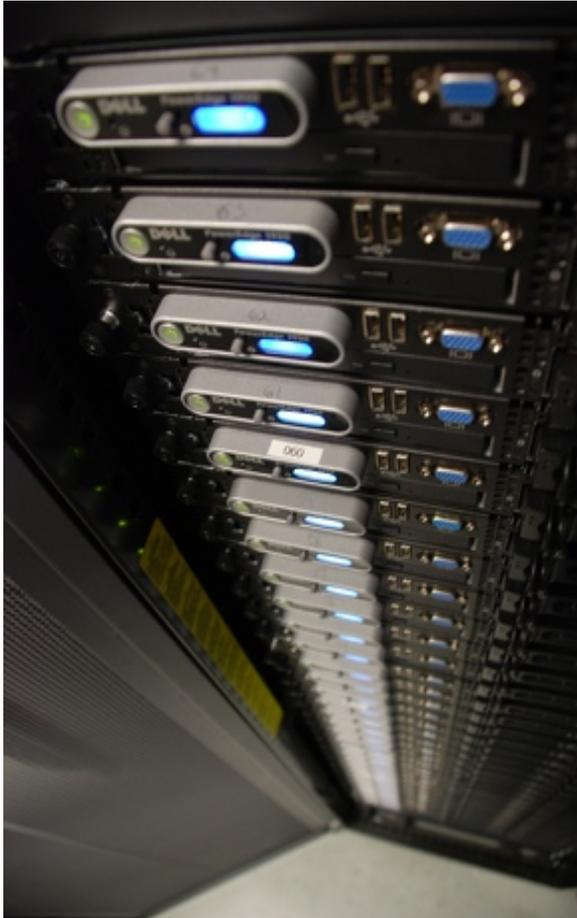


With the lights off...





Front and back of Painter



The high-bandwidth, low-latency InfiniBand interconnect, along with the LONI optical network, is ideal for **moving large datasets in and out of LONI clusters** enhancing local U.S. ATLAS data analysis.



Building a Local Tier 3 Center with Grid Access (T3g)

TIER 3 COMPUTING FACILITY



New Tier 3 with Grid Access (T3g)

- With the awarded DOE ARRA Stimulus funds, we will *soon* purchase computing hardware to build a scalable, robust and dedicated Tier 3 analysis center.
- The new Tier 3 cluster will allow us to retrieve datasets from the grid, submit pAthena jobs to the grid, and run local analysis jobs.
 - Local HEP physicists and students will have **exclusive access** – no outside analysis/production jobs from the grid
 - Perfect for **quick analysis runs** where fast turnarounds are required
 - Longer and more demanding analysis jobs will be done on the LONI clusters
 - Tier 3 storage will be for **current and recent datasets** and shared with LONI clusters through a SE. The PetaShare data grid and tape system will be used for **long-term storage** of analysis data and results.
- A future expansion to a Tier 3 center with Grid Services (T3gs) is probable considering the extensive amount of computing and storage resources we have around the state of Louisiana.



Planned Tier 3 Hardware

- **Compute nodes:**

- Dell PowerEdge R410/R510 rack-based servers
 - Intel Xeon 5600 series, 6-cores per socket (12-cores per node)
 - 24 GB of memory per node to best utilize the triple-channel memory bus (6 x 4GB)

- **Dedicated storage:**

- Depending on the configuration, we will use either Lustre, Hadoop HDFS, or Xrootd.
 - Dell PowerVault MD1200 direct-attached storage (12 SATA 6Gb/s)
 - Or, using worker node storage (R410: 4 disks, R510: 8-12 disks)

- **Cluster and grid servers:**

- Dell PowerEdge R410/R510/R710 servers
- Virtualized cluster head nodes and grid services, particularly a BeStMan SRM/ GridFTP Storage Element (SE) sharing data with LONI clusters



Computing Facilities at Louisiana Tech

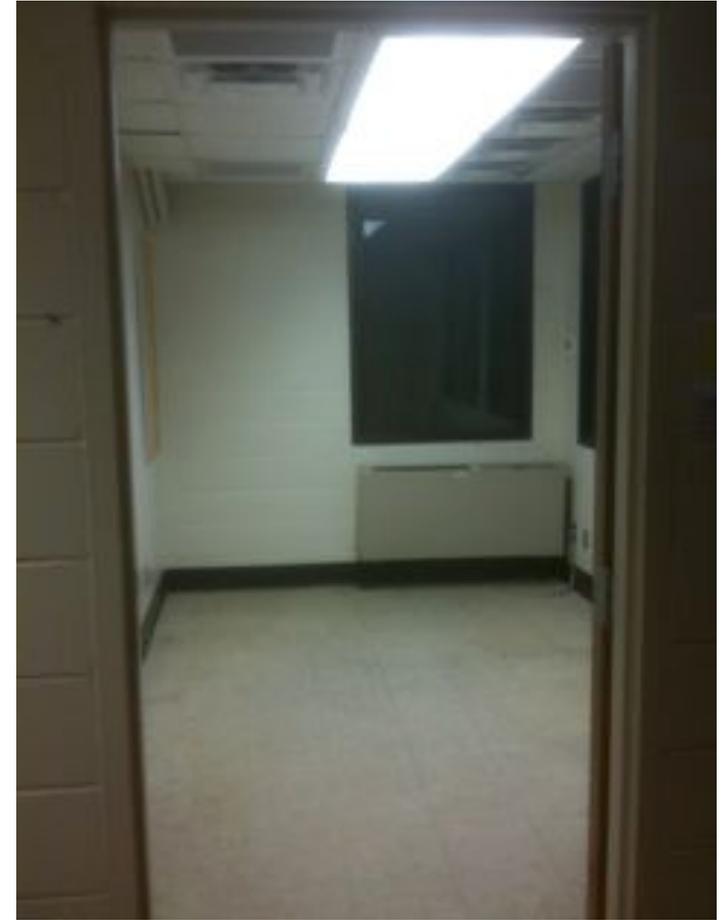
- Engineering Annex Data Center **New!**
 - Small computing facility located at the Physics department
 - **Four rack spaces** available (room size: 8' x 14')
 - 12.5 tons of cooling (150,000 BTU/h)
 - Supporting close to **44 kW of power**
 - Funded by University Research and the Center for Secure Cyberspace
 - Temporary home for their virtualization cluster (3 racks) for about 2 years
 - Our new **Tier 3g cluster** will also be located here
- Data Replication Center
 - Houses two of the LONI clusters on campus
 - **Painter** (Linux, 512 cores)
 - **Bluedawg** (IBM AIX)
 - Mostly used as backup facility for Louisiana state agencies
 - Extra spaces are available for LaTech clusters, if needed



Tier 3 Computing Facility



The currently empty computing facility located at the Physics department in the Engineering Annex building. The room was completed in March 2010.



Plenty of space, power, and cooling for four rack spaces. It's not much, but it is better than what we had with the old A/C unit in back of the room.



OSG Compute Element (CE) and DZero MC Production

GRID COMPUTING EFFORTS



OSG Compute Elements

LONI_OSG1 (active)

- Official LONI CE (osg1.loni.org)
 - Located at LSU
- OSG 1.2 production site
- Managed by LONI/LaTech staff
- Connected to “Eric” cluster
 - Opportunistic PBS queue
 - **64 cores** out of 512 cores for *serial jobs* (also shared with other PBS queues)
- Primarily used to support DZero production

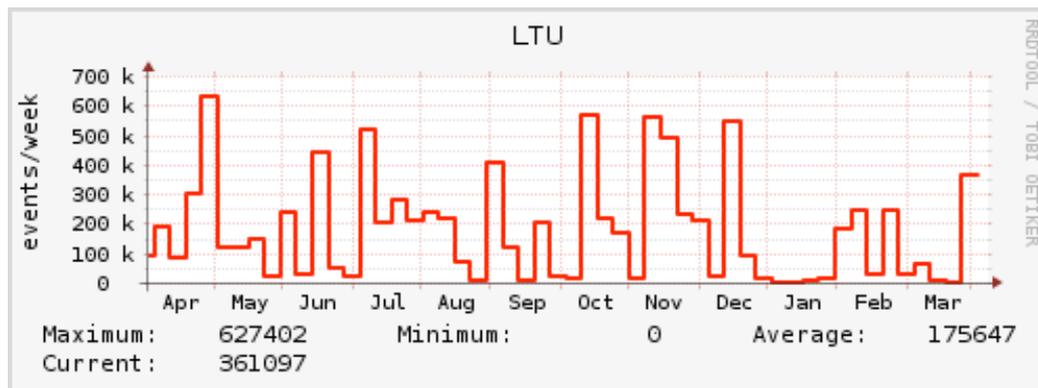
LTU_Tier3 (*planned*)

- Tier 3 CE (ce1.grid.latech.edu)
 - Located at Louisiana Tech
- Enable LaTech researchers to submit analysis jobs from the grid back to our T3g
 - Particularly useful for LaTech researchers located at CERN
- No outside analysis jobs from the grid would be allowed



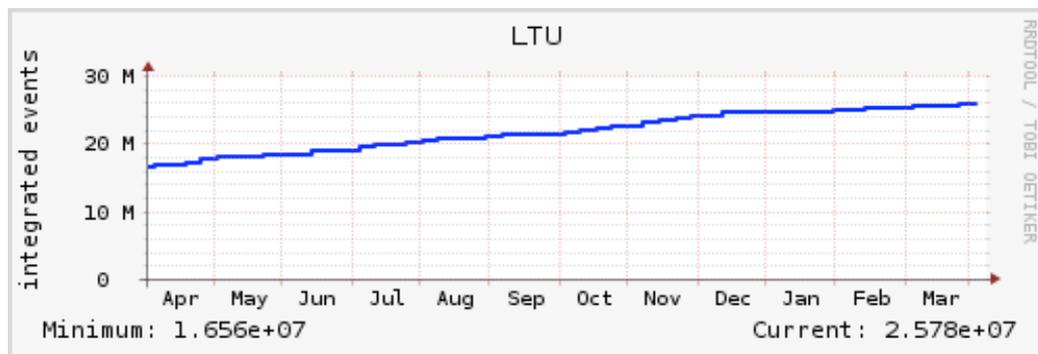
Current Status of LONI_OSG1

- Continuing to run DZero MC production on LONI cluster at LSU



Weekly DZero Monte Carlo production

MC production varies as Physics requests become available



Cumulative DZero Monte Carlo production

LONI_OSG1 has produced **~9 million** Monte Carlo events in the last year and **~25.8 million** events in the past six years on our small clusters



Closing Remarks

- Our HEP researchers have been exploring the limitations of our local Tier 3 and LONI clusters for local analysis.
 - Both faculty and graduate students have been using ROOT/DQ2/Athena for a while now on our local Tier 3.
 - The CPU and disk limitations of our Tier 3 led to the DOE ARRA Stimulus funds.
- LONI and PetaShare provide us with innovative computing and storage solutions that fully exploit the high-bandwidth, low-latency optical network technologies of LONI.
 - We plan to use the six LONI Linux clusters for our longer, demanding analysis jobs. Smaller jobs will be done on our new Tier 3 cluster (less wait time).
- We'll continue to pursue ATLAS analysis on LONI using the 10 TBs of storage from PetaShare.
 - A common storage area for ATLAS software, such as Athena, needs to be established on LONI clusters before major adoption, though.
- We look forward to building a local Tier 3 Center with Grid Access (T3g) soon.



QUESTIONS / COMMENTS?