



# LTU Site Report

Dick Greenwood  
Louisiana Tech University

April 5, 2007

# LTU & CCT

## Computing Resources



	CPU(GHz)	Storage (TB)	People
LTU	54.8 GHz	1.28	1F+1PD+1GA+1UG
CCT (Helix)	128*2* 2 GHz	0.1	
LONI (in 2007)	(5*132 + 720)*2.33	>250 x10 <sup>3</sup>	

People:

Dick Greenwood: Institutional Rep.

Joe Steele: DOSAR representative in data reprocessing group, Certifying data file merging for all of D0's remote reprocessing farms.

Michael Bryant: system manager; OSG installations manager, High Availability computing research (with Box Leangsuksun)

Vishal Rampure : High Availability Computing research (also with Box)

John Rollo: student system manager in training

# Personnel Using LaTech IAC (CAPS) Linux Cluster



- **At present:**
  - **HEP Group (D0, ATLAS)**
    - Lee Sawyer, Dick Greenwood (IAC Rep.), Markus Wobisch, Joe Steele
  - **Jefferson Lab Group (G0, Qweak Experiments)**
    - Kathleen Johnston, Neven Simicevic, Steve Wells
- **Other users**
  - **Box Leangsuksun, LaTech CS**



# CAPS Linux Cluster Overview

**Windy Cluster**

- 23 total CPUs
- 54.8GHz total
- 16GB of RAM
- 1.28TB of storage

Debian 3.1 (Sarge, i386)  
OSG-enabled (soon) cluster



**Caps10: SAMGrid head node and SAM station**

- Pentium4 2.0GHz
- 1GB of RAM
- 440GB of RAID1 storage

gigabit network connectivity



**Windy0 and Windy1**

- Dual Intel Xeon 2.2GHz with Hyper-Threading
- 2GB of RAM
- 140GB of storage



**Windy2 and Windy3**

- Dual Intel Xeon 2.8GHz with Hyper-Threading
- 2GB of RAM
- 120GB of storage



**Windy4 and Windy5**

- Dual AMD Opteron 2.0GHz
- 2GB of RAM
- 80GB of storage



**Windy6 and Windy7**

- Intel Pentium4 2.4GHz
- 2GB of RAM
- 80GB of storage

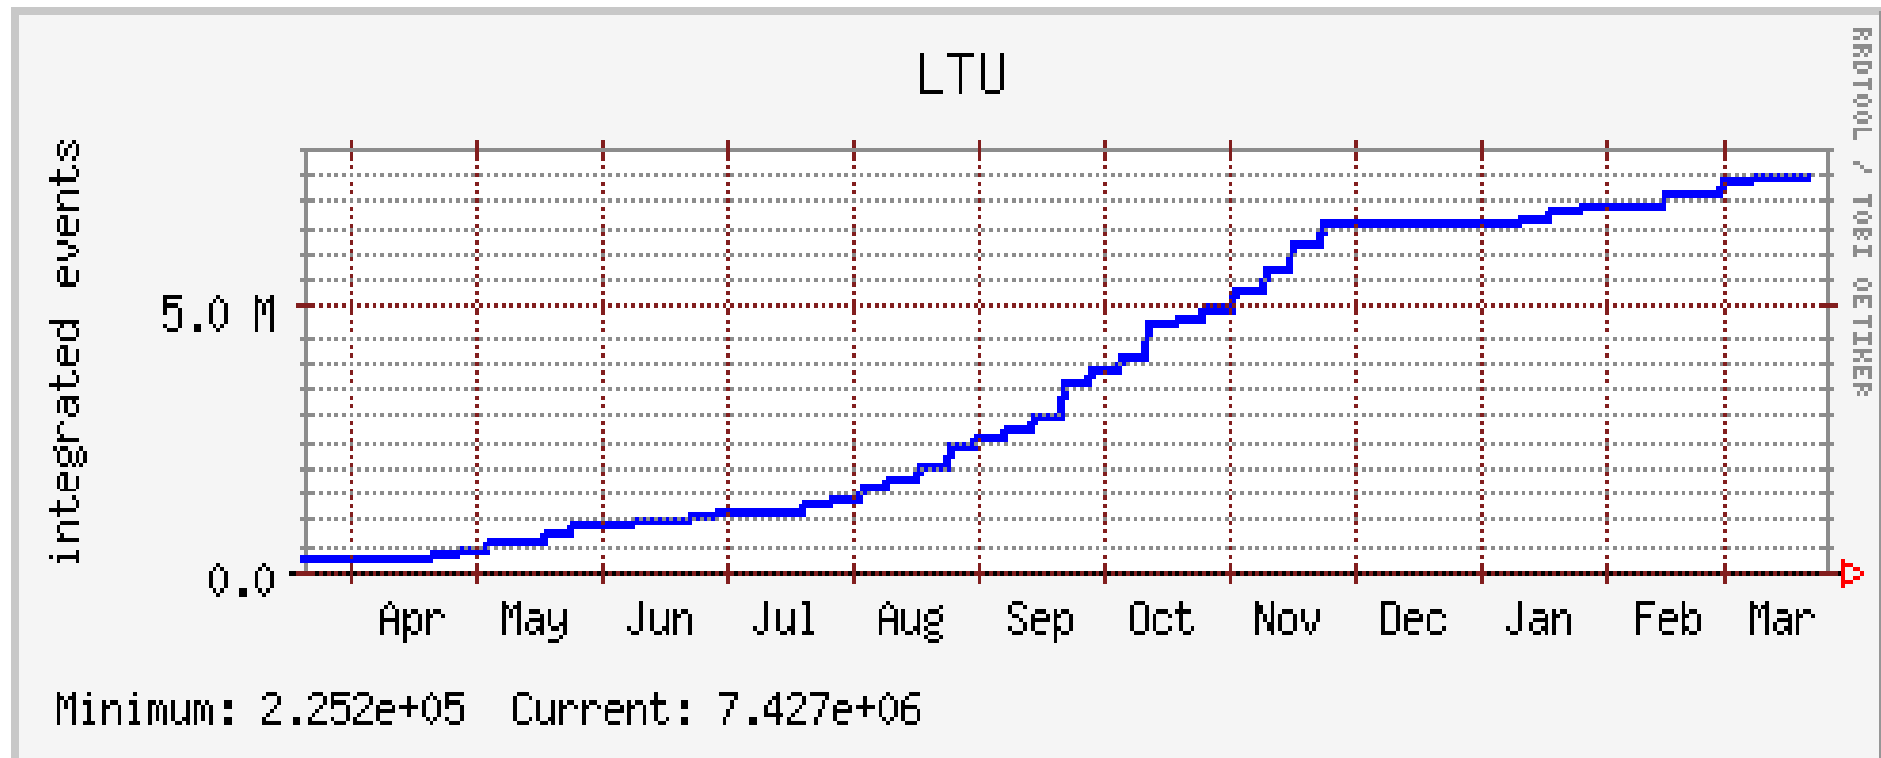
# Status of CAPS Site



- **Locally at LTU:**
  - Our local SAMGrid head node and SAM station, caps10, has been updated to match the installed products and configuration of the LSU CCT head node, Itu-cct.
- **Creation of LTU\_OSG:**
  - The OSG worker node client software and Condor is installed on caps10.
  - The OSG CE is installed on caps10 and early in 2007, became registered as OSG Production site, the first in the state!
  - Some MC production has occurred but production is currently halted until the completion of the LONI installation to the LaTech campus.



# From the D0 MC Production Page





## **SAMGrid Execution Site: Itu-cct**

### **Sam Station: Itu.cct.Isu.edu**

- Operating as SAMGrid site since Dec. 2005
- Certified to run P17.09.06 with Sam V7
- Currently running jobs submitted by J. Snow
- In 2007, installed OSG software: now a registered D0 and DOSAR Production site:
  - D0 MC Production and Remote Reprocessing site
  - Available for future ATLAS Production

# Louisiana Optical Network Initiative (LONI)



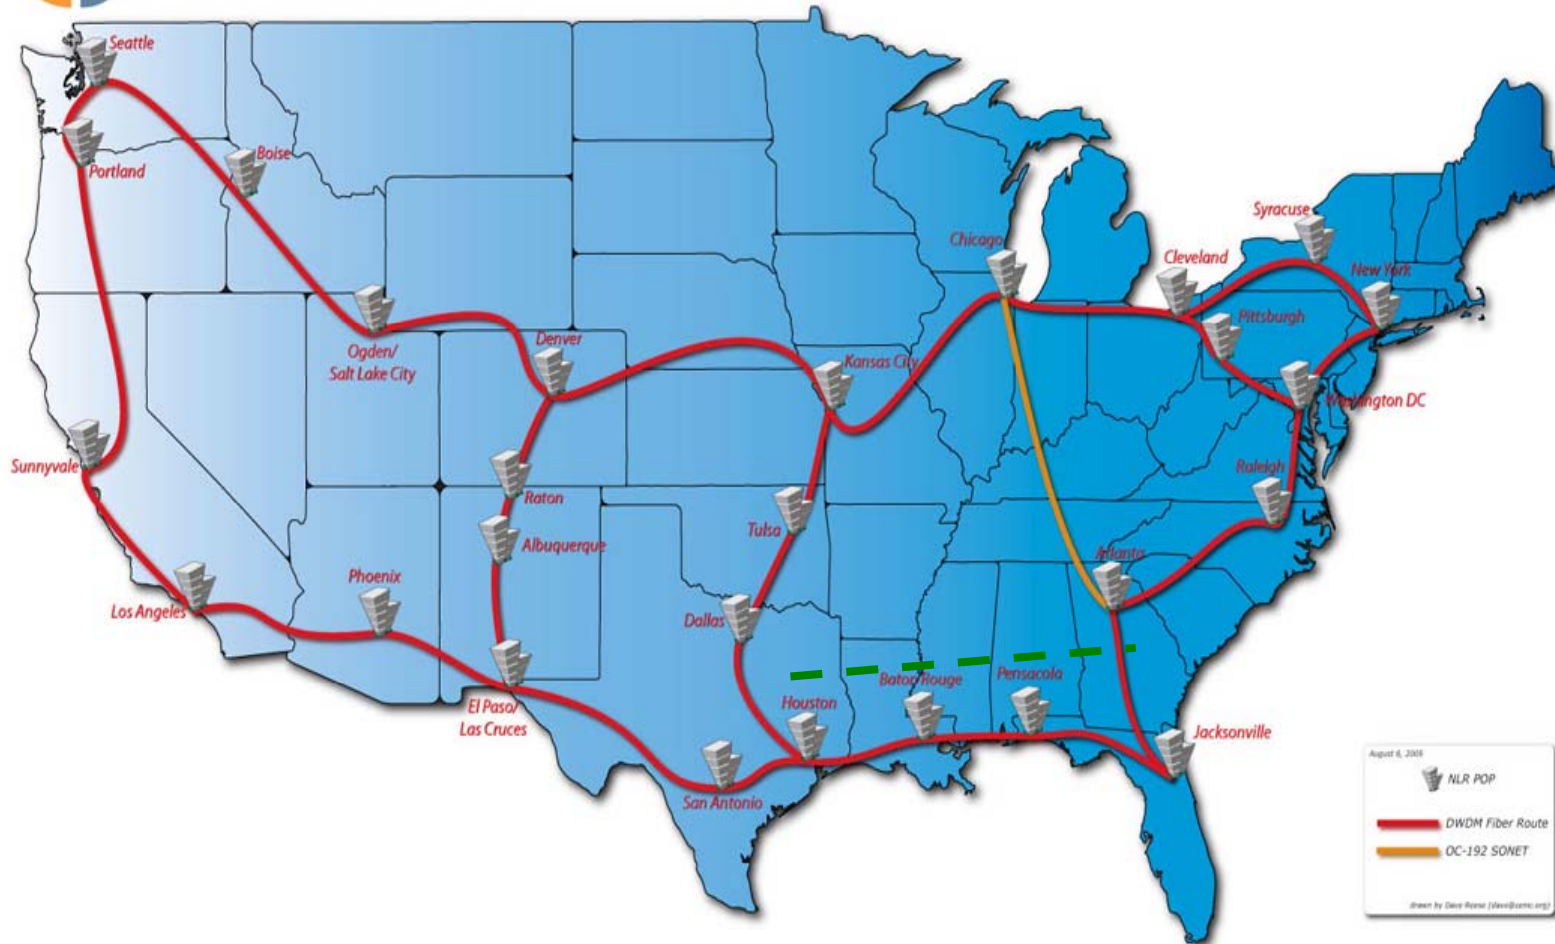
- In September 2004, the Louisiana State has committed \$40M for a state-wide optical network.
- 40Gb/sec bandwidth
- Spanning 6 Universities and 2 Health Centers:
  - LSU
  - Latech
  - UL-Lafayette
  - Tulane
  - UNO
  - Southern University
  - LSU Health Centers in
    - New Orleans
    - Shreveport







# National LambdaRail™ Infrastructure



© 2005 National LambdaRail™

For more information regarding NLR see <http://www.nlr.net> or contact [info@nlr.net](mailto:info@nlr.net)

— 2001 Plan



# New Dell Linux Clusters



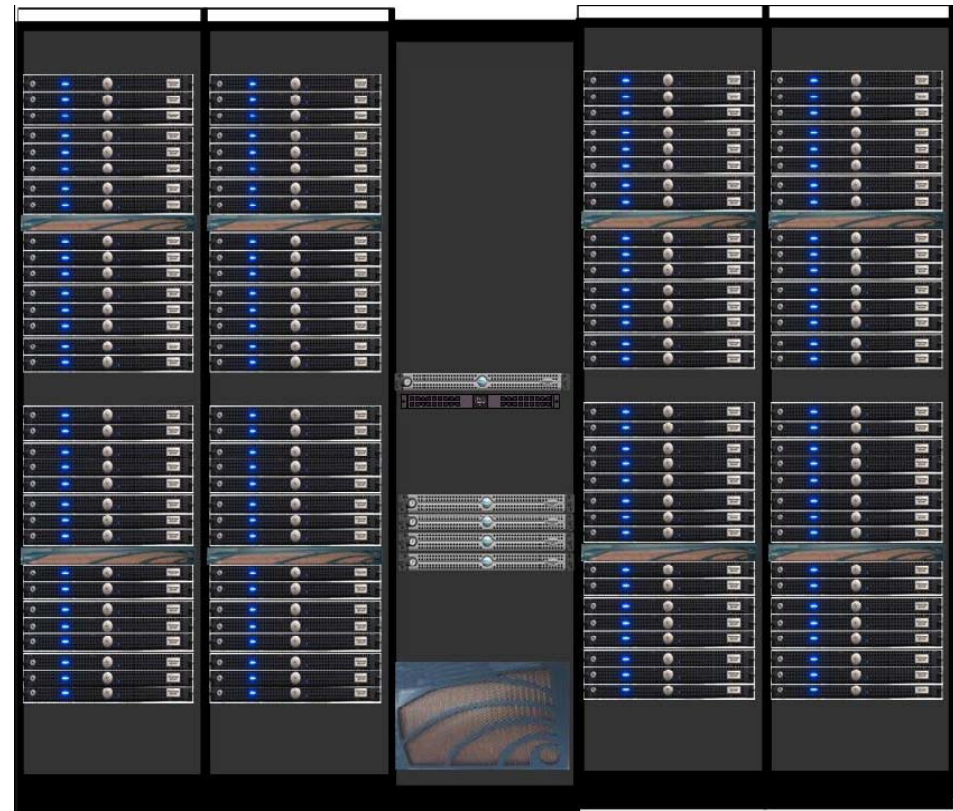
- “Remote” Systems
  - 132 nodes, four-way Intel “Woodcrest”  
2.33 GHz, 4 GB RAM [4.921 TF; 528 GB RAM]
  - First deployment at LSU; 384 processor’s dedicated to D0 Reprocessing Mar-Apr, 2007
  - Next deployment at UNO; deployment at Tech this summer (2007).
  
- Central System in Baton Rouge
  - 720 nodes, eight-way Intel “Cloverto(w)n”  
2.33 GHz, 4 GB RAM [53.683 TF; 2.88 TB RAM]
  - To be deployed soon.

# Dell Cluster Description



## Environmentals:

- 208 V / 310 amperes; 64.5 KW; 18 tons of cooling (7,868 cfm)
- 6 racks total (4 node, 1 control, 1 storage)
- Rack Dimensions:  
78.7"Hx23.94"Wx  
39.93"D
- Each rack has 4 PDUs;  
4 L6-30 208v connects  
(total of 24 L6-30  
circuits)

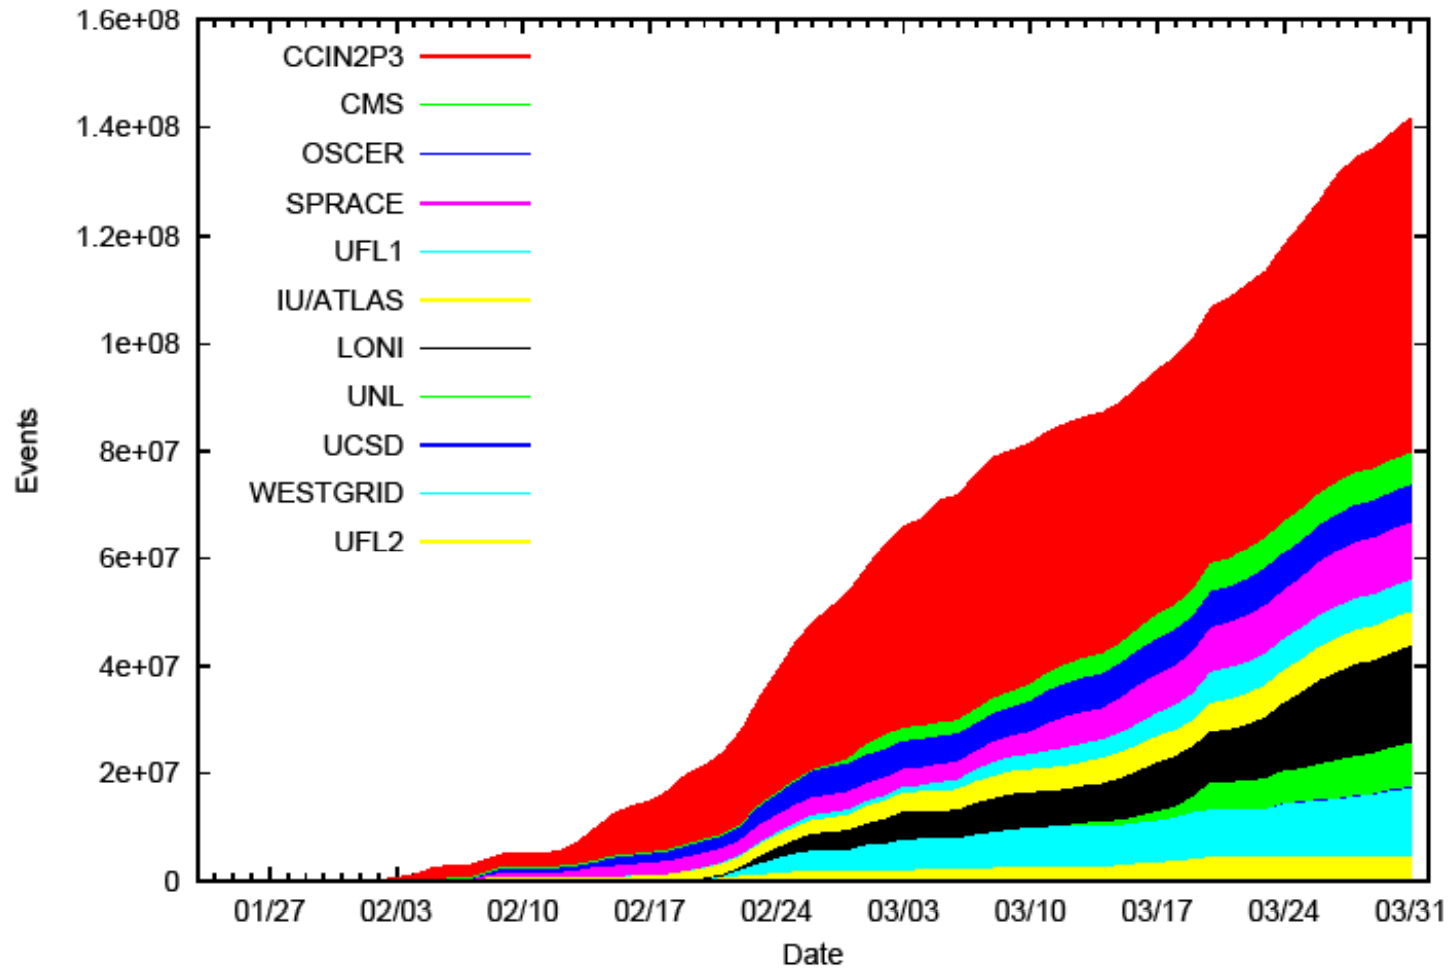


# Schedule for Deployment of LONI Dells



- All dates are based on fuzzy power & air conditioning availability projections by the sites which keep getting pushed back
- LSU 5 TF: to be completed in two weeks
- UNO 5 TF            April 16
- ULL 5 TF            April 30
- ISB 50 TF            May 15
- LaTech 5 TF        June 4
- Southern 5 TF      June 18
- Tulane 5 TF         July 2
- LONI-wide OSG CE established: June 1

Remote p20 Production to 01-Apr-2007





# Storage



Currently very tight:

- /home is 25 GB and /scratch is 280 GB
- this limits usability
- all served via Network File System (NFS) which is not high performance
- Future:
  - When central Linux cluster comes, it will include:
    - 14.0 TB raw at each “remote” site in one rack
    - 225 TB raw at the central site
  - Will provide central /home storage as well as global /scratch space
  - Using Lustre filesystem supported by Cluster Filesystems, Inc.

# Award of NSF MRI: PetaShare



1 x Overland Storage NEO Series 8000 tape storage system -  
400 TB Capacity: \$ 153,558

Specifications : Model 500, 12 Native Fibre LTO-3 Tape Drives, 500 Tape Slots, Redundant Power, Remote Access & Control, Touch Screen Front Panel (Quote includes shipping and assembly)

Deployment site: Louisiana State University


5 x Overland Storage REO Series 9000 disk storage system - 44 TB capacity  
each: 5 x \$96,144 = \$480,720

Specifications : 44TB Raw SATA, 38 TB usable RAID 5, Protecton OS,  
Remote Access and Control,  
Rackmount (Quote includes shipping and assembly)


Deployment sites: Louisiana State University, Louisiana Tech University,  
Tulane University, University of Louisiana at Lafayette, Tulane  
University, University of New Orleans

Total requested equipment cost: \$634,278









- **Goal:** enable domain scientists to focus on their primary research problem, assured that the underlying infrastructure will manage the low-level data handling issues.
- **Novel approach:** treat data storage resources and the tasks related to data access as first class entities just like computational resources and compute tasks.
- **Key technologies** being developed: data-aware storage systems, data-aware schedulers (i.e. Stork), and cross-domain meta-data scheme.
- **Provides** an additional 200TB disk, and 400TB tape storage



Participating institutions in the PetaShare project, connected through LONI. Sample research of the participating researchers pictured (i.e. biomechanics by Kodiyalam & Wischusen, tangible interaction by Ullmer, coastal studies by Walker, and molecular biology by Bishop).



# LTU Plans



- Participate in the development of the Louisiana Optical Network Initiative (LONI) with its 10GB/sec connectivity to the National Lambda Rail and T2
- Continue association with DOSAR and the OSG in Grid Computing applications and research
- In collaboration with CCT and LONI, will incorporate large storage and new Dell cpu's within the next few months
- Continue production work for D0
- Commence collaborating with SW Tier 2 and with the CCT and HEP program at LSU on ATLAS analysis



# Extra Slides

# Tevfik's Schedule for PetaShare



Year 1: The required equipment for the development of PetaShare will be purchased, installed at each participating site, calibrated, and will be integrated with the existing equipment at these sites. In parallel to this process, we will start developing the key technologies necessary for this instrumentation: data-aware storage, data-aware schedulers, and remote data analysis and visualization.

Year 2: Technologies developed during the first year will be integrated. Transparent storage selection, data archival, cataloging and remote access techniques will be developed on top of this integrated system. A user friendly and uniform interface will be developed for interaction.

Year 3: The first prototypes will be deployed at the participating sites. Testing and further development will be performed. Application groups will start using the new instrumentation actively.

Year 4: The developed system will be ported to other platforms. It will be made publicly available to the community. Commercialization potential will be investigated.

We are currently in the process of negotiating with different vendors. Status = ?

# PetaShare Overview


LaTech

ULL

LSU

UNO

Tulane



5 x IBM P5  
w/ 112 proc

1.2 TB RAM

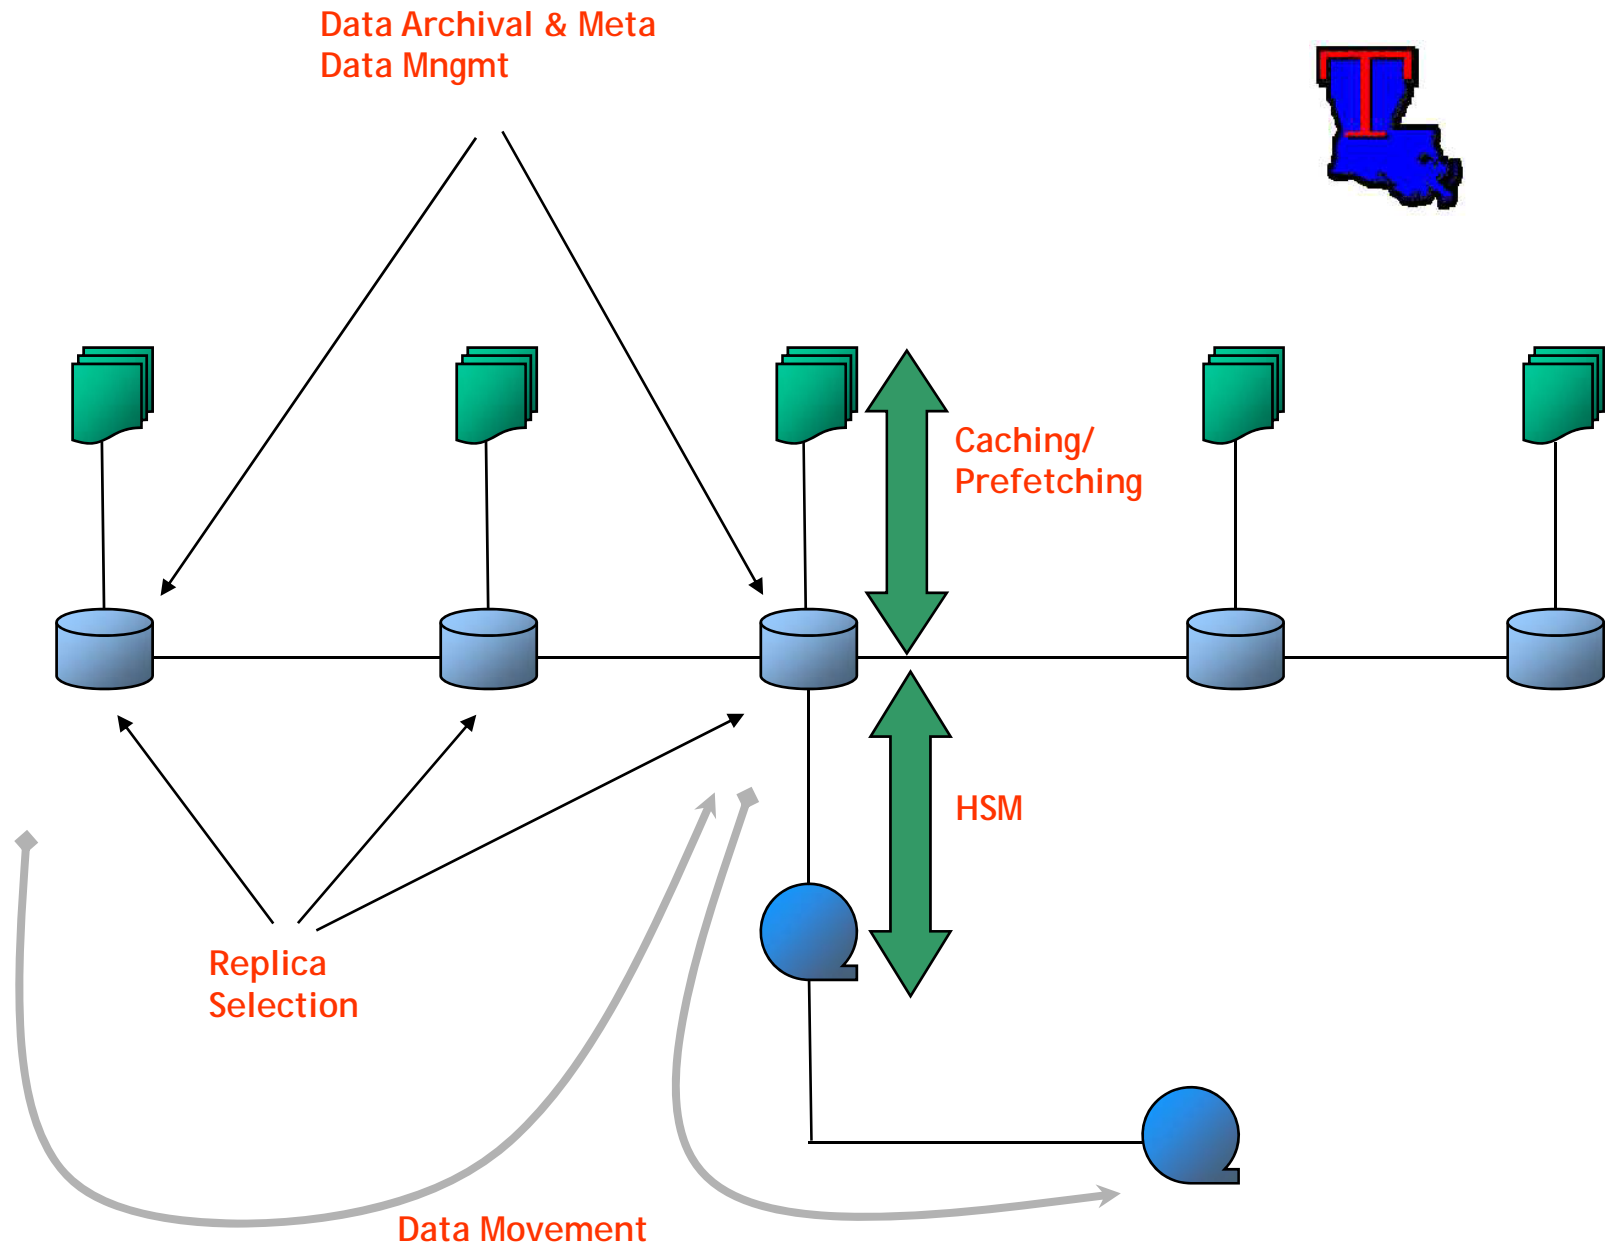
200 TB Disk

400 TB Tape

SDSC

*Transparent  
Data  
Movement*

5APR07





# People involved with PetaShare Development and Usage

	Senior Personnel			Postdocs			Graduate Students			Undergraduates		
	Women	Minority	Total	Women	Minority	Total	Women	Minority	Total	Women	Minority	Total
LaTech			2						5			2
LSU	3		27			5	5	4	47	2	2	24
Tulane	1		3			2	3	5	11		1	3
ULL	7	4	18			8	13	5	31			4
UNO			7				5	4	32	3	8	17
Total	11	4	57			15	26	18	126	5	11	50