

# Site report: Wisconsin

Sridhara Dasu

Dan Bradley, Will Maier, Ajit Mohapatra  
{dasu,dan,wcmaier,ajit}@hep.wisc.edu

University of Wisconsin - High Energy Physics

USCMS T2 Workshop - FNAL, 2010.03.08



- 1 Overview**
- 2 2009-2010 Facilities Status**
  - Software
  - Batch
  - Storage
  - Network and Facilities
- 3 2010 Deployment Plans**
  - Batch, storage and software
- 4 Storage development and Plans**
  - Storage tools
  - dCache is dead, long live dCache?
- 5 Analysis at Wisconsin**



Component	Software
OS	Scientific Linux 5.3
Batch	Condor 7.4.0
Storage	dCache 1.9.5-8
Grid	VDT 2.0.0p10
Data transfer	PhEDEx 3.2.9
Frontier	Squid 2.7.STABLE7

**Table:** 2010 Software Status

- Software stack increasingly stable from batch to storage
- Emphasis on deployment, monitoring
- Relocatable OSG installs would (still) be nice




CPU Class	HS06	Slots
2 x 2.80 GHz Xeon	840	86
2 x 3.00 GHz Xeon	125	10
4 x 1.80 GHz Opteron	1710	240
8 x 2.66 GHz Xeon	1976	120
8 x 3.00 GHz Xeon	2229	260
Dedicated	6880	716
Opportunistic	-	2952

**Table:** 2010 Batch Status

Since last workshop:

- Added more opportunistic resources (~4000 HS06)
- Worked to minimize efficiency cost of preemption policies on opportunistic resources
- More work on Condor scaling, glide-ins (Dan Bradley)





**Figure:** Dedicated and opportunistic slots, 2009-2010



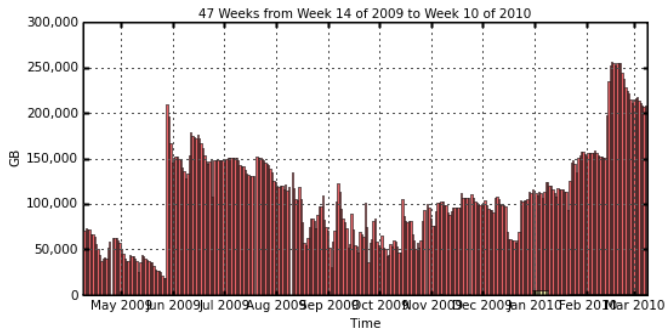
Dedicated	280 TB
Dual-purpose	200 TB
Raw	480 TB
Writable	240 TB

**Table:** 2010 Storage Status

Since last workshop:

- Purchased 480 TB of dedicated storage; burning in now (online next week)
- Upgraded PNFS server hardware, faster RAID (200% performance improvement)
- Developed near real-time dCache replicator (billingrep)
  - No file loss since Summer 2009





■ GLOW


■ GLOW-ATLAS

■ GLOW-CMS

Maximum: 255,876 GB, Minimum: 0.00 GB, Average: 107,147 GB, Current: 207,806 GB

**Figure:** dCache usage, 2009-2010





**Figure:** Campus WAN usage, 2009-2010


- Set new campus records
- Worked with campus to plan better peering to Nebraska/i2
- Added more power to old machine room
- Planning for improved cooling in new machine room





## CMS PhEDEx - Transfer Rate

24 Hours from 2009-12-15 16:00 to 2009-12-16 16:00 UTC



Maximum: 761.39 MB/s, Minimum: 0.00 MB/s, Average: 164.46 MB/s, Current: 3.21 MB/s

**Figure:** PhEDEx transfers during peak network usage, 2009



Month	Storage (raw/writable TB)	Slots	HS06
March	480/240	-	-
June	256/128	256	2400
July	Chimera?	-	-
Current	480/240	716	6880
New in 2010	720/360	256	2400
Total in 2010	1200/600	972	9280

**Table:** 2010 Hardware Deployment Plan

- Already close to 2010 milestones
- Dedicated storage: 8 x 2.66 GHz Xeon, 8 GB RAM, 24 x 2 TB disk
- Dual-purpose: 8 x 2.26 GHz Xeon (75 HS06), 24 GB RAM, 4 x 2-3 TB disk



- Developed many tools to manage dCache
  - All code available at <http://code.hep.wisc.edu/dcache-tools/>
  - Includes `dcache_billingrep` (replication service), monitoring scripts
- Developed generic Python package to support command line tools
  - Documentation and code at <http://packages.python.org/pyCLI/>
  - Used in `dcache-tools`; intended for reuse elsewhere
  - Provides argument/config parsing, logging, daemonizing, unit/functional test support, etc
  - (`dcache-tools` uses it)



- 2009: Participated in early look at Hadoop
  - Hardware failures and dCache bugs forced us to switch focus back
- Since Summer 2009, storage has been reliable
  - Few failures (mostly hardware), but they're well understood
  - dCache can handle the load at our site
  - billingrep and other local tools worked around many problems
- ...but Hadoop still offers a few things that we really like:
  - Chunking
  - Native replication
  - Rack awareness
- Now that dCache is mostly stable, we have an opportunity to take another look
  - Plan to resurrect small Hadoop test stand Summer 2010
  - Still need to move to Chimera
  - Long term? Whatever works



- Affiliated with: FWD, TSG (EWK)
- Positive feedback from groups, regular use of site for analysis
- October Exercise (2009)
  - EWK Electron group produced skims at Wisconsin and published data to other T2s
  - Hosted other Electron data for analysis at Wisconsin; no problems
- Quickly find batch problems when supporting global MC production (Ajit Mohapatra)
- farmout scripts support local analysis (Dan Bradley)

