

# BNL Site Report

HEPIX @ TRIUMF

Tony Wong

Oct. 25, 2011

# General Facility News

## ● Power & Cooling

- Average ~810 KW usage as of September 2011
- Usable capacity of ~1.8 MW of UPS-backed power
- Cooling capacity finally caught (surpassed?) heat load since 2005 crisis
- Option to grow to ~2.8 MW of UPS-backed power (not needed yet)
- Frequent maintenance issues with power & cooling equipment

## ● Mass Storage

- Considering new (7<sup>th</sup> overall) robotic silo for RHIC (early 2012?)
  - Increase in RHIC capacity by 10,500 LTO-4 tapes
- Added 10 more LTO-5 tape drives (on top of 20 purchased earlier) to RHIC
- Week-long, major upgrade of HPSS (version 7.3.2) in Dec. 2011
  - Migrating core from AIX-based servers to RHEL servers
  - ACSLS (library control) platform migration from AIX to Solaris x86
  - DB2 storage upgrade

# General Facility News

- New Linux Farm additions
  - Purchased ~29 kHSPEC for US-ATLAS in June
    - Total of 848 servers with ~100 kHSPEC of computing power
    - About ~1.9 PB of local unmanaged storage capacity
  - Purchased ~27 kHSPEC for RHIC in September
    - Total of 924 servers with ~100 kHSPEC of computing power
    - About ~5.6 PB of local managed storage capacity
  - Purchased ~4 kHSPEC for other programs in 2011
    - Total of 87 servers with ~9 kHSPEC of computing power
    - About ~500 TB of local managed storage capacity
- Continue to experience some QA problems with Dell hardware
  - Memory issues (frequent single-bit errors) in 2011 cluster
  - Overheating issues (loose heat sinks) with 2010 cluster

# Overheating CPU's (first reported at Cornell Fall 2010)



no thermal grease



with thermal grease

# General Facility News

- **Storage Management - ATLAS**
  - Local Tier 3 Xrootd deployment now part of US Federation (in addition to Tier 1 dCache) with global redirector at BNL
    - Interface to WAN via proxy server
    - Implemented File Resource Manager (FRM) with backend script that queries global redirector and can pull missing files from Tier 1
    - DQ2-get integrated with xprep to pre-populate Xrootd storage
    - Implemented checksum verification with checksum passed from client
- **Xrootd monitoring via Ganglia using scripts developed at Dubna**
- **Overall US Federation plans still under discussion**
- **Storage Management – RHIC**
  - Added ~70 TB to PHENIX dCache scratch pool
  - dCache upgrade for PHENIX pending

# General Facility News

- NFS

- Built a RH-based NFS 4.1 client for testing, but...
- Beta-test of BlueArc NFS 4.1 postponed till mid-2012
- No luck with Netapp either
- Apparently no market-driving force to support NFS 4.1 anytime soon, but discussions with both BlueArc and Netapp are on-going

- Other news

- Migration from Webauth web SSO to Shibboleth web SSO for authenticated pages almost complete (still need to migrate certain internal and limited visibility websites)
- On going migration of all US ATLAS network equipment to CDCE (new data center built in 2009)
- Implementing dedicated “IP block” for infrastructure/grid servers

# Unix Centralization Project

- Mandated by funding agency
  - Affects all computational devices (desktops, servers, printers, etc)
  - Standardized configuration management tools
  - Automatic updates for vulnerability patches
  - Centralized administrative privileges
- Update since Cornell HEPiX meeting
  - Chose Puppet for configuration management (**completed**)
  - Updated RACF security compliance document (**completed**)
  - Migrate data center servers to Puppet management (ongoing)
  - Migrate staff Unix desktops to Puppet management (ongoing)

# Experimental Support

- Support for “small” experiments
  - Neutrino experiments (Daya Bay, LBNE)
  - Astrophysics (LSST, DES, BOSS)
  - AnDY (new RHIC experiment)
  - RHIC R&D (EIC)
- Disproportional support effort compared to established experiments (RHIC and ATLAS)
  - Tendency to “reinvent” the wheel
  - Will receive one extra staff to support 7 (and counting) “small” experiments

# Software News

- Software

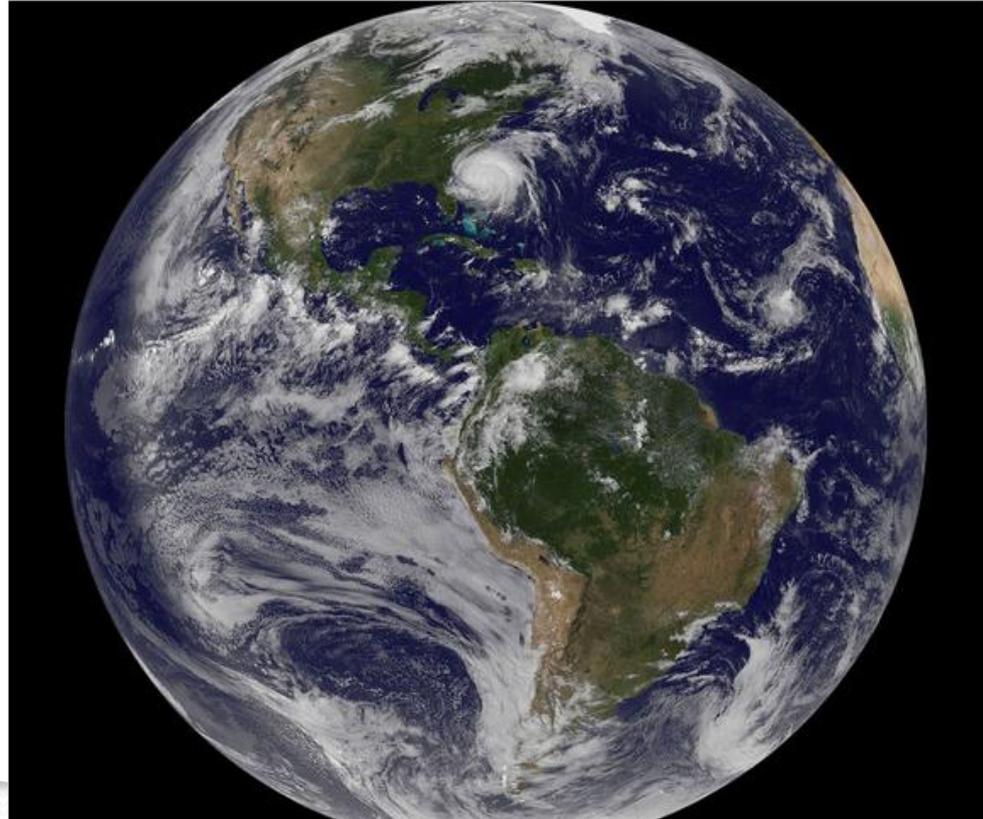
- Upgrade to Condor 7.6.4
  - Late 2011 until early 2012
  - Improved support for virtual environment and cloud computing
  - More flexible configuration management for group quotas
- Scientific Linux 6.x
  - Currently using SL 5.3
  - Waiting for SL 6.2
- Ksplice
  - RACF has 5-year annually renewable contract
  - Used at RACF for reboot-less updates of Linux-based systems
  - Important for security updates on critical systems
  - Bought out by Oracle – no changes in support structure so far

# Other Facility Developments

- Cloud computing
  - 32-node testbed
  - Deployed CERNVMFS and Condor
- 10 Gb connectivity to Linux Farm
  - Tied to evolution of multi-core technology
  - Evaluation of Cloud Edge server with Interlagos (AMD) chip
    - Dual motherboard capable of supporting 4 cpu's (16-core) each
- Migration to Chimera
  - 2-day outage for distributed storage cluster (ATLAS)
  - Replaces PNFS as dCache name service
  - Nov 2011
- Four staff openings due to staff turn-over

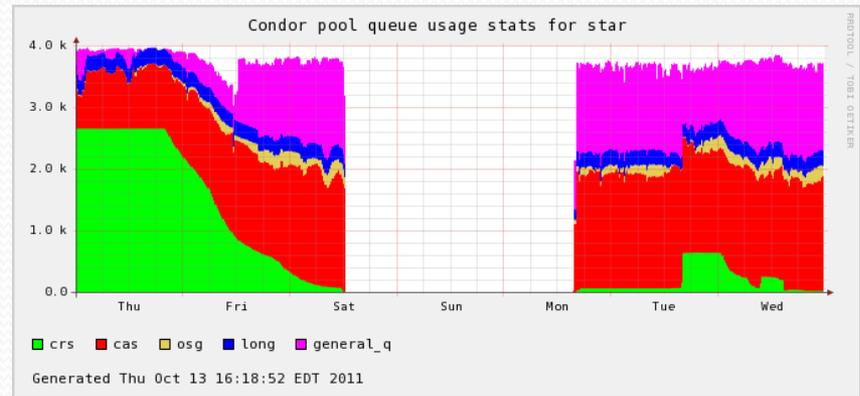
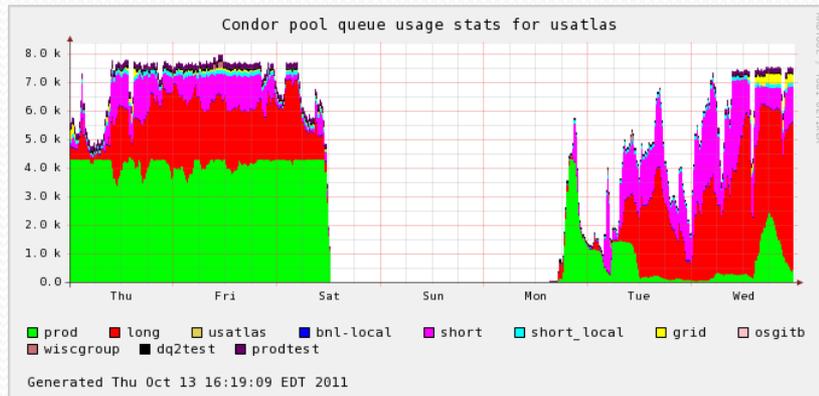
## Hurricane Irene

- Facility shut down for precautionary reasons on Saturday, August 27 at 1 pm.
- Facility restart on Monday, August 29 at 12 noon.
- Luckily, eye of hurricane passed over Manhattan, 90 km to our west
- Brookhaven did not lose electrical power during storm
- About 50% of homes on Long Island lost electrical power, and it took ~one week for full restoration



# Aftermath of hurricane shutdown

- Many failed disks in PHENIX distributed storage
  - Not surprising – disks ridden hard in PHENIX computing model
  - 48 dead disks (262 TB of storage capacity) shortly after power up following hurricane – almost all of them older (3+ years) systems
  - Another 3 disks failed within the next 3 weeks (tail of distribution)
  - Worker node-based distributed storage
- More statistics
  - About 3 hours to complete shutdown of RACF
  - About 5 hours to start up all RACF services



# Organizing CHEP



- Nearly on schedule—according to internal Gantt chart
- Abstract submission deadline this Friday (Oct. 28)
- Registration opens in mid-November
- 2<sup>nd</sup> Bulletin to be published before registration opens
- Lot of work
  - Help from NYC event organizing company
  - Contracts for all venues (meetings and social events) signed
  - Thanks to Alan Silverman for handy document (Guide To Organizing CHEP)!
  - More info at [www.chep2012.org](http://www.chep2012.org)