BNL RHIC/ATLAS Computing Facility Site Report

Christopher Hollowell <hollowec@bnl.gov>

RHIC/ATLAS Computing Facility (RACF)

Physics Department

Brookhaven National Laboratory

RHIC/ATLAS Computing Facility (RACF) Overview

Created in the mid 1990s to provide centralized computing services for the RHIC experiments: BRAHMS, PHOBOS, STAR and PHENIX

Expanded our role in the late 1990's to act as the tier1 computing center for ATLAS in the United States

Now also supporting some smaller computing facility installations for LBNE, Daya Bay, EIC and LSST

Currently employ 31 FTEs

RACF Overview (Cont.)

- LHC shutdown for the next 2 years for upgrades
  - ATLAS not collecting new data, but continuing large-scale processing of previously amassed data
- RHIC Run 13 began in February
  - Started with polarized proton collisions
  - Switching to Gold-Gold
  - Plan to continue run until June 2013
  - Expect ~2 PB (~1 PB STAR, ~1 PB PHENIX) of new raw data

Datacenter Infrastructure

- 14,000 square foot datacenter
  - See our business continuity talk later this week for more info
- Power Redundancy
  - 1 MW battery UPS
  - 1.3 MW diesel generator, with 2 flywheel UPS systems
    - Have had some flywheel and generator malfunctions
      - Critical hosts now dual-powered via utility and generator
- Cooling
  - Basement AC
  - CRAC units
  - Liebert XD Overhead units
    - Problems with R134a leaks - phasing out in favor of CRACs

Datacenter Infrastructure (Cont.)

- Switching to metered APC floor PDUs with remote monitoring
- New Server Technology CDUs for some racks capable of remotely powering on or off individual power outlets
  - Necessary for equipment without out of band management capabilities
- Synapsense
  - Being used to monitor temperature, humidity, underfloor air pressure, and CRAC temperature/humidity in two rooms
    - Adding monitoring to the rest of the datacenter
      - ~500 sensepoints currently, moving to ~2500

Datacenter Infrastructure (Cont.)

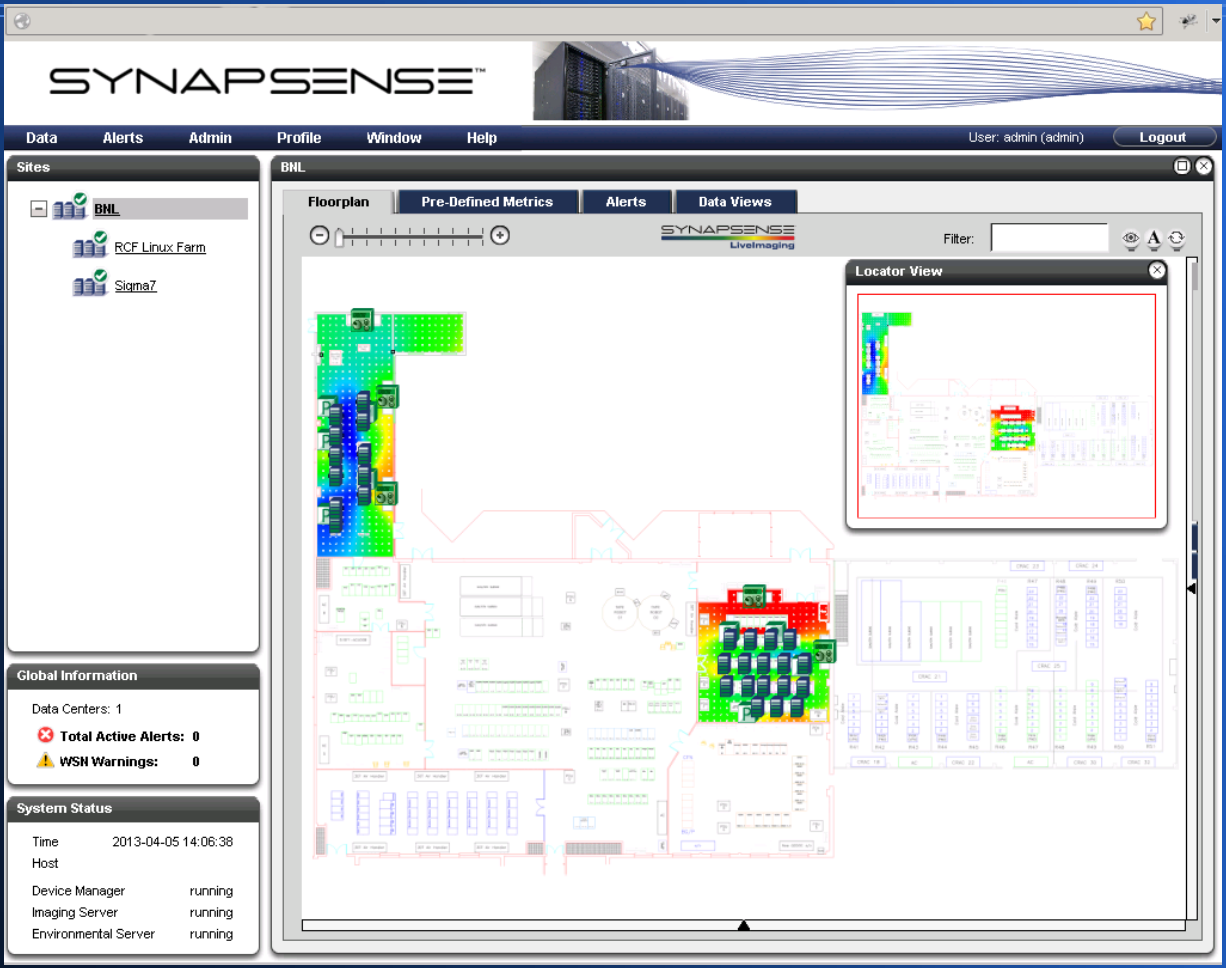

Synapsense Web GUI

Mass Storage

- Using HPSS as our backend mass storage system
  - Upgraded to 7.3.3 p6 in December 2012
- ~31 PB of data currently in tape
  - ~42,000 tapes
- 7 Oracle/StorageTek SL8500 tape libraries
  - 6 in production, 1 empty
- 2 StorageTek Powderhorn 9310 silos
  - Planned on retiring, but couldn't due to repacking overhead
- ~310 TB total disk cache for HPSS
  - IBM DS3400 and DS3500 arrays

Mass Storage (Cont.)

BNL developed tape scheduler for file retrieval - ERADAT[1]

Using LTO5 tape for RHIC Run 12/13 data

All HPSS servers, including core, running Red Hat Enterprise Linux (RHEL) 5

- Won't move to RHEL6 until the next release of HPSS – likely next year

HPSS mover network upgraded

- 40 GigE capable
- Now a fully independent network
- Network connectivity to counting houses now redundant fiber

ATLAS STK SL8500 Tape Libraries

NFS Storage

- Primarily using BlueArc appliances for NFS service
  - Mainly for user home directories and scratch space

- 3 BlueArc clusters serving ~975 TB usable storage

- RHIC – 6 Titan 3200 heads, split into 2 clusters
  - 19 LSI arrays; 2 HUS-150 arrays; 1 HUS-130 array
  - 2x10 GigE connections per head
  - ~908 TB usable storage

- ATLAS – 2 Mercury 100 heads
  - 4 LSI arrays
  - 2x10 GigE connections per head
  - ~70 TB usable storage

NFS Storage (Cont.) – Recent Upgrades

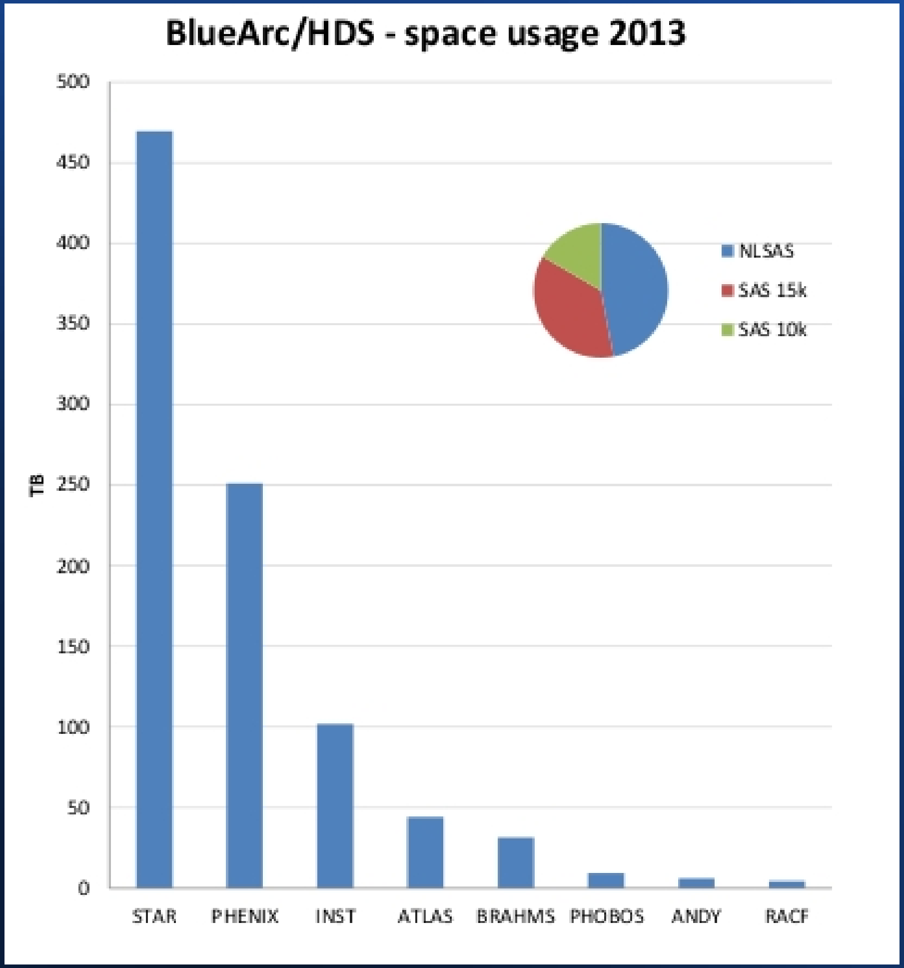

Increased STAR capacity by 15%

Retired 7 old PHENIX Thumpers with high cost of maintenance and poor performance

- Migrated to BlueArc
- Increase of 255% in storage for PHENIX

New PHENIX storage consists of 288 2.5” 10K 900 GB SAS disks behind 2 HUS-150 controllers

AFS Storage

- Upgraded to OpenAFS 1.6.1 on RHIC AFS fileservers ~6 months ago
  - RHIC fileservers running as RHEL6 VMs under Red Hat Enterprise Virtualization (RHEV)
  - 1.4.14/RHEL5 still on all DB servers, and USATLAS fileservers

- Primarily used for experiment software repositories

- Using Teradactyl TiBS for backups

- RHIC - ~5 TB of space served by 4 file servers
- ATLAS - ~3 TB of space served by 2 file servers

Distributed Storage

- dCache
  - ATLAS
    - Version 1.9.12 with Chimera DB on FusionIO
    - ~12 PB total disk space
  - PHENIX
    - Upgrading to 2.2.x with Chimera
    - ~3.2 PB total disk space
      - Most of this storage capacity is provided by the local disk in the processor farm nodes (~550)
      - The majority of the data is also in tape

Distributed Storage (Cont.)

- XRootD
  - STAR
    - ~3.4 PB storage space on ~500 processor farm nodes
    - 2 redundant redirectors
    - Supervisors spread across different racks, and floor PDU feeds to enhance resiliency
  - ATLAS
    - BNL Local Tier3 – ~130 TB storage
    - Participation in ATLAS FAX (Federating ATLAS Storage systems using XRootD))
      - Operate US global redirector and XRootD interface to Tier1 dCache storage

Network

Considering moving to connectivity beyond 1 Gbps for processor farm nodes

- What are others doing?
  - Channel bonding, Infiniband (IP over IB), staying with 1 GigE attached nodes, 10 GigE?
  - Total cost of ownership (TCO)?
  - Speak with me offline

Effect of PHENIX Job Ramp-Up on Traffic: 31.5 GB/s (252 Gbps) max

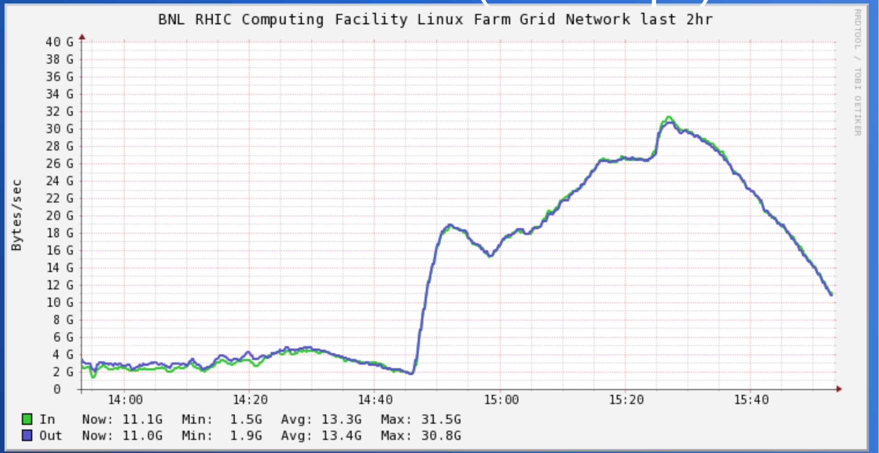

Looking at 100 GigE WAN-wide connectivity as part of BNL science DMZ

- Currently channelized 10 GigE lines (~80 Gbps)

Looking at interswitch links beyond LACP 10 GigE: currently ~160 Gbps

General Services

- WWW servers, DB servers, SSH/SFTP gateways, centralized monitoring, DNS servers, mail servers, LDAP servers, testbeds, etc.
  - Currently ~100 machines and ~300 virtual hosts
    - ~100 VMs managed via libvirt
  - Many high availability services migrated to RHEV
    - Using version 3.0
    - 12 systems in RHIC RHEV cluster, 10 for ATLAS
      - ~100 VMs in each cluster

- All hosts running RHEL5 or RHEL6: system deployment and management via Cobbler, Puppet and Red Hat Network (RHN) Satellite

Cloud Activities

- BNL RACF Private Cloud
  - OpenStack Essex release; 32 nodes supporting ~400 VMs
  - VMs decoupled from direct access to RACF network storage (NFS, dCache DCAP, etc.)
    - *lcg-cp* used for data transfers in ATLAS jobs
    - CVMFS used for ATLAS software releases
  - EC2 interface

- BNL_CLOUD Panda Site/Queue
  - Maps to VMs running in the RACF private cloud, as well as Amazon EC2
    - Using EC2 spot pricing – VMs evicted if cost rises too high
  - Also experimenting with Google Compute Engine
  - VMs created using Boxgrinder

Cloud Activities (Cont.)

- BNL_CLOUD Panda Site/Queue (Cont.)
  - VMs submitted via condor_g (EC2 interface)
    - AutoPyFactory (APF) maintains a static VM target count
      - Working on dynamically submitting VMs based on demand
    - VMs run condor_startd daemons which join a pool managed by Condor central manager systems at BNL
      - Condor-C with shared password authentication
      - Using Condor connection broker (CCB)
    - Scalability issues with this setup overcome
      - 3,000+ condor_startd <--> 1 condor_collector communication over WAN overloaded the collector
        - Now running 20 collectors aggregated via CondorView

Cloud Activities (Cont.)

- BNL_CLOUD (Cont.)
  - Ran well for ~3 weeks with 5,000+ simultaneous simulation jobs during testing in January

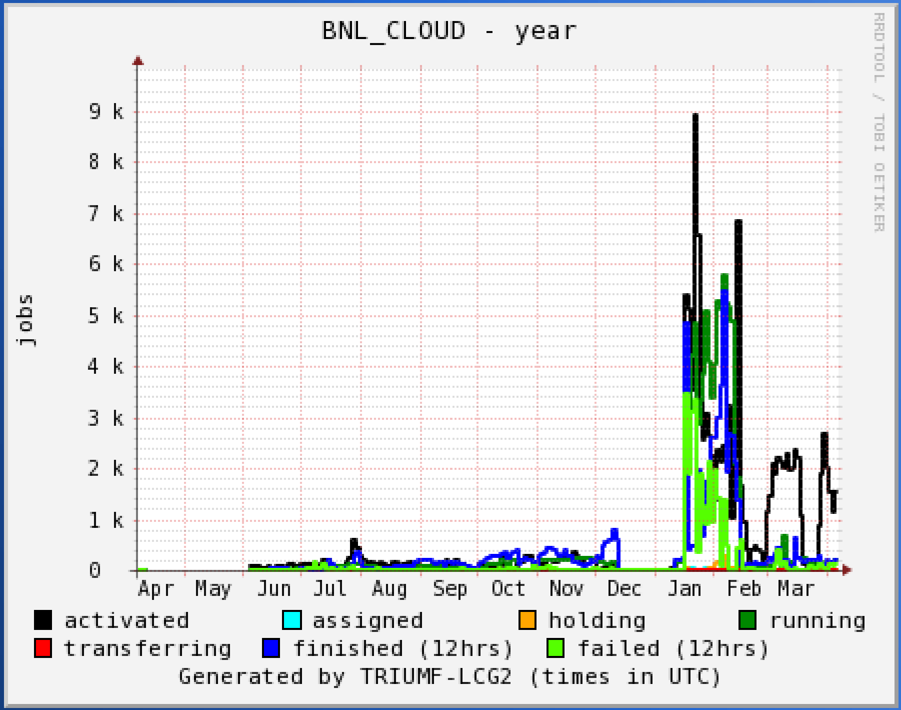

- ATLAS High Level Trigger (HLT) Farm
  - ATLAS wanted to make use of its HLT farm to run simulation while LHC is offline
    - Installed OpenStack Folsom CERNVM-based images
  - BNL RACF sent a system administrator to CERN for a few weeks to help with a test setup
    - Additional involvement anticipated

Processor Farm

~2,000 systems, providing ~32,000 logical CPU cores

Purchased 182 Dell R410 hosts for RHIC in September 2012
- Dual 2.8 GHz Xeon X5660 processors
- 48 GB DDR3-1333 RAM, 4x2TB SATA drives
- 4,368 additional logical cores
- Last round of Westmere based systems we'll purchase

Retired the last of our Penguin Computing servers in the fall
- Only Dell machines in our processor farm at this time

Bid out for 90 new Sandy Bridge based ATLAS servers

Successfully tested SL6 UEFI PXE and booting off of >2TB drives

Processor Farm (Cont.)

New RHIC Dell R410 servers
with Liebert XD Overhead Cooling

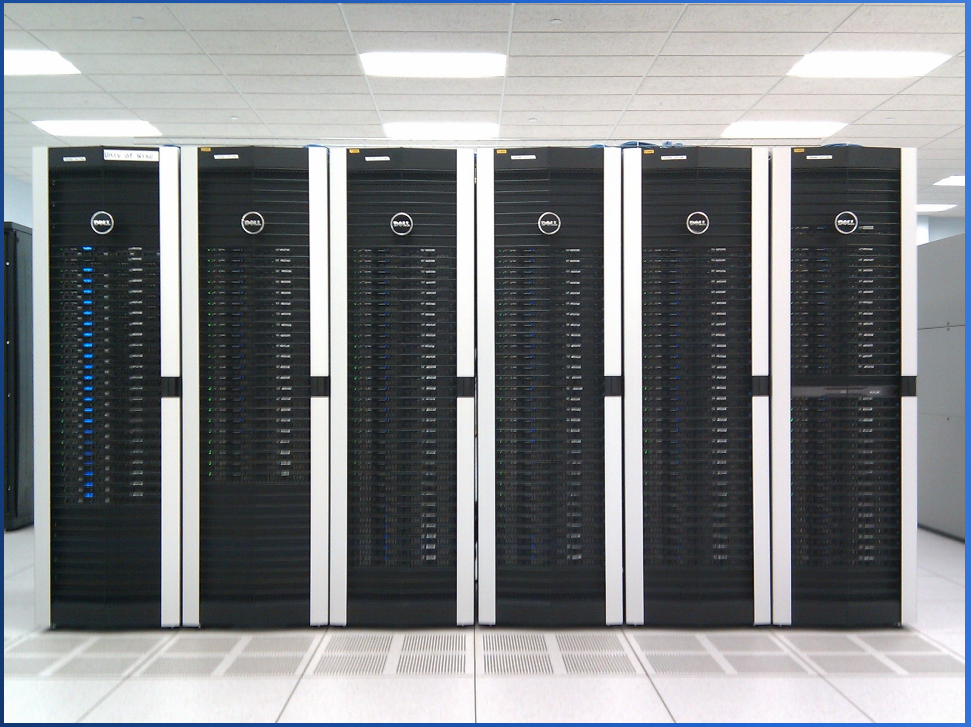

ATLAS Dell R410 Servers

Processor Farm (Cont.)

- Systems currently running 64-bit SL5
  - Upgrade to SL6 planned in the near future
    - ATLAS: expect to upgrade in the next two months
    - RHIC: late summer, or early fall 2013

- Condor used as our batch system
  - RHIC: 7.6.4
  - ATLAS: 7.6.6
  - Planning on upgrading to the latest stable version available during the farm SL6 upgrade
  - Problems using hierarchical group quotas with dynamically partitioned slots
    - Can't enable the GROUP_ACCEPT_SURPLUS feature

Processor Farm (Cont.)

- Condor (Cont.)
  - Beginning to allow users to instantiate VMs as jobs
  - Not using Condor's VM universe because it can't restrict images run: designed our own system and image set. See CHEP2012 paper for more information[2]
    - Being used by STAR to run efficiency calculations (simulation) in SL4 VMs
      - Related to data production from 2008
      - Code release had been “frozen” - STAR required the same environment for comparison of results

- Investigating use of SSDs on processor farm hosts
  - Price per GB improving but still expensive
  - Also started looking into hybrid SSDs via FlashCache & Bcache

Processor Farm (Cont.)

Considered using SSDs for local scratch storage. However, random write performance (with Crucial CT512M4SSD2) wasn't much better than a SATA (Hitachi HUA723020ALA640) drive tested. Still useful for random read intensive applications

Multi-threaded (24) aggregate, buffering disabled, bonnie++ -b -r 2560 -s 5120

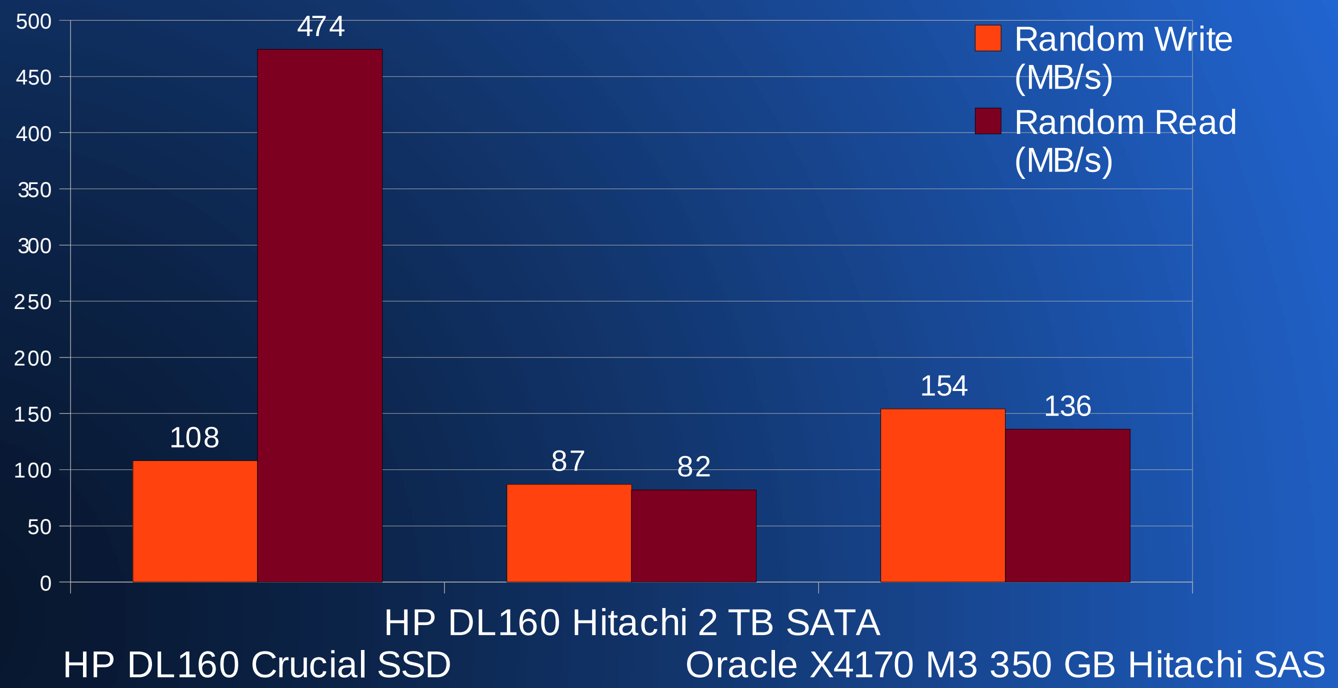

Sandy Bridge CPU Benchmarking

~30% increase in HEPSPEC06 performance for 32 logical core Sandy Bridge CPUs when compared to “equivalent” class Westmere CPUs

Systems Tested

1. Dell R420: 2 E5-2470 CPUs, 64 GB DDR3-1600 RAM
2. Dell R720xd: 2 E5-2660 CPUs, 64 GB DDR3-1600 RAM
3. Oracle X4170 M3: 2 E5-2690 CPUs, 64 GB DDR3-1600 RAM
4. HP ProLiant DL160: 2 E5-2630 CPUs, 64 GB DDR3-1333 RAM
5. Dell R410: 2 X5660 CPUs (Westmere), 48 GB DDR3-1333 RAM

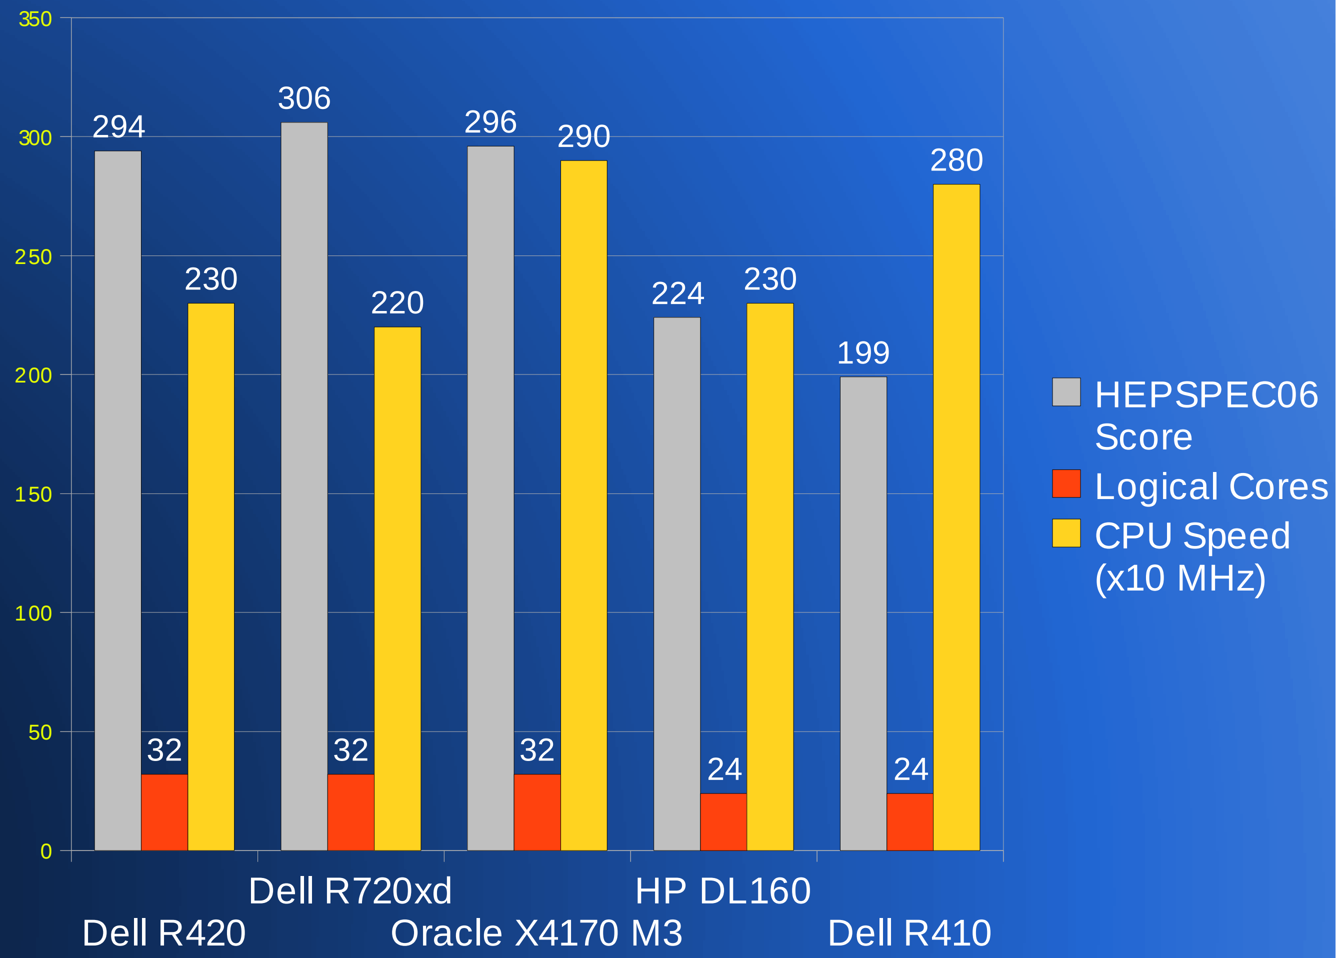

Sandy Bridge CPU Benchmarking (Cont.)

Power footprint appears largely unchanged by Sandy Bridge

Power utilization of servers based on mid-range Sandy Bridge CPUs is similar or lower when compared to those based on mid-range Westmere CPUs

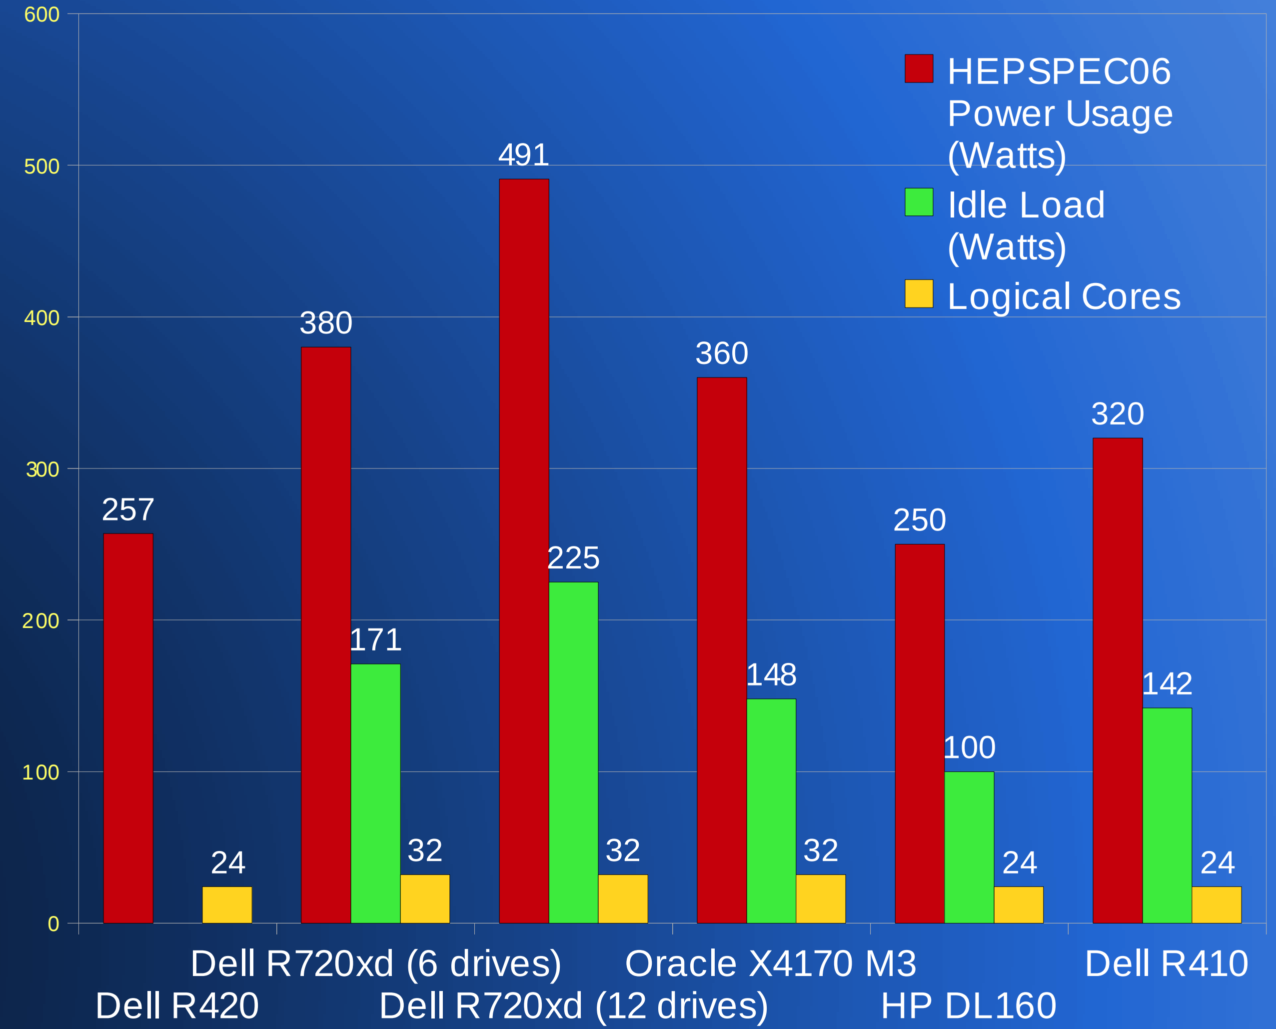

Ksplice/Oracle Uptrack

- Allows one to patch a running Linux kernel without rebooting
- In use on our processor farm, and critical infrastructure hosts
- We've been using Ksplice/Oracle Uptrack for over 2 years
  - Very happy with the software
  - No issues or crashes/panics encountered: updates are of a high quality
- Significantly reduced administrative time needed to deal with kernel patches, and has minimized vulnerability exposure
  - Scheduling rolling reboots for the farm was necessary before
- Updates released quickly, sometimes even before they are available for SL
- We have SL/RHEL support because we were an existing customer
  - Oracle Linux is the only enterprise distribution supported for new customers

Ksplice/Oracle Uptrack (Cont.)

“Raw” Ksplice utilities still available and usable to create/insert one's own rebootless kernel patches on any Linux distribution

- Released under the GNU GPL
- http://oss.oracle.com/ksplice/software/

Works by inserting a module containing functions modified by a patch, and replacing the first instruction of each modified function in kernel memory with a *jmp* instruction to the corresponding patched function

- First necessary to quiesce the kernel via *stop_machine()*

Caching/service of Uptrack updates available via a local server

Web GUI allows one to see the status of all of one's hosts

See my CHEP 2012 paper for more information[3]

Ksplice/Oracle Uptrack (Cont.)

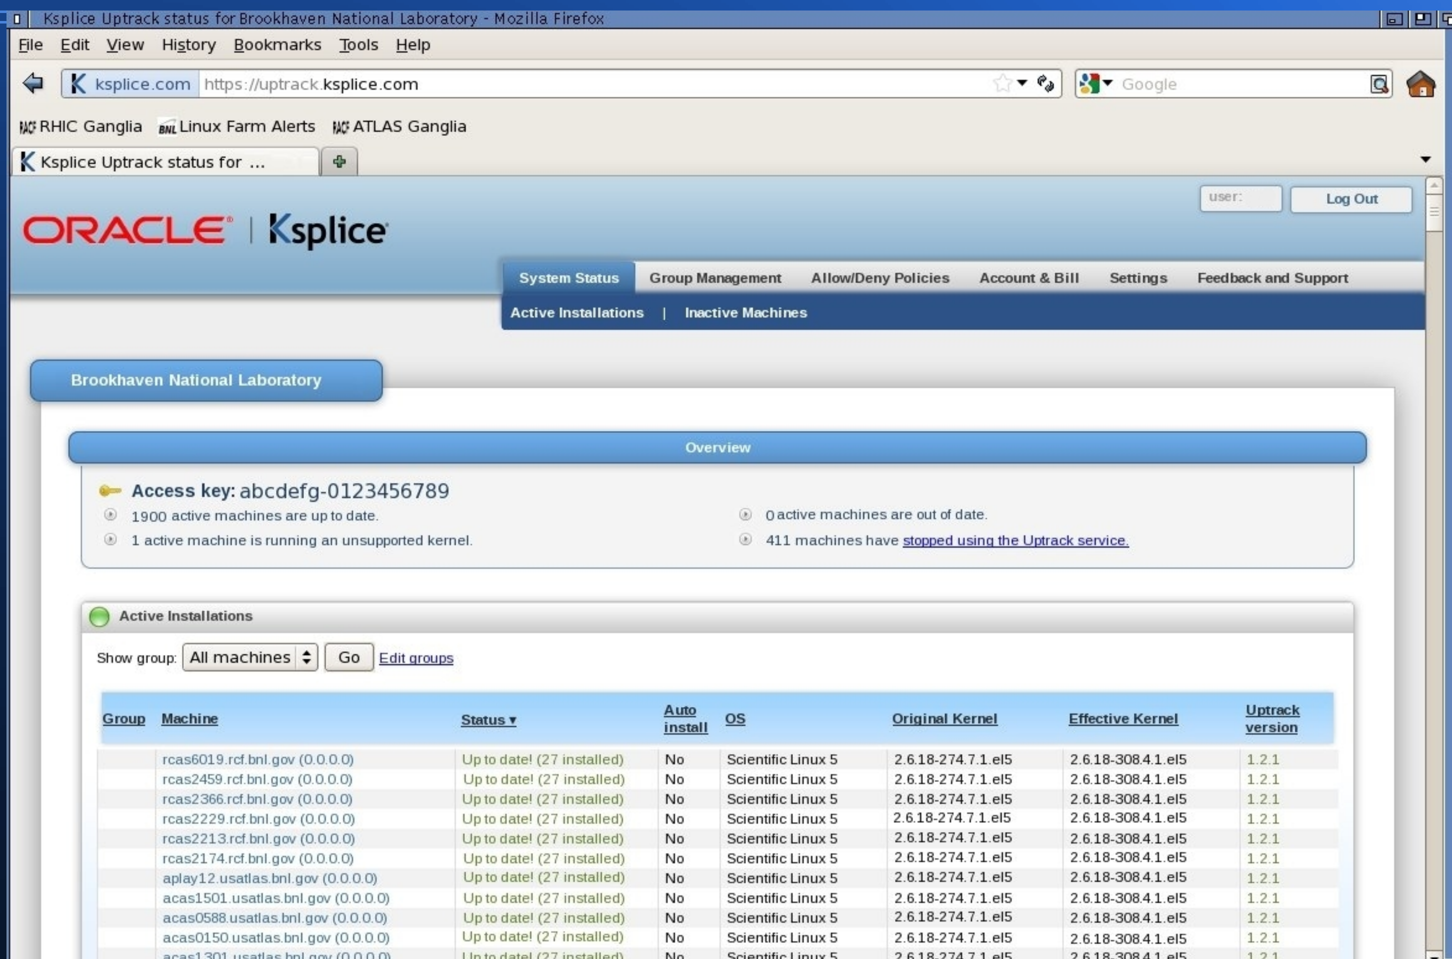

Ksplice Uptrack Web Status GUI

Questions?

Thanks to the following people at BNL for contributing some of the information presented:

Costin Caramarcu, Richard Hogue, John Hover, Hiro Ito, Jerome Lauret, John McCarthy, Shigeki Misawa, James Pryor, Tejas Rao, Ofer Rind, Jason Smith, Will Strecker-Kellogg, Tony Wong, David Yu, Alex Zaytsev, Xin Zhao

References

1. Tape Storage Optimization at BNL
D Yu, J Lauret - Journal of Physics: Conference Series, 2011
http://iopscience.iop.org/1742-6596/331/4/042045

2. Simplified Virtualization in a HEP/NP Environment with Condor
W Strecker-Kellogg et al -Journal of Physics: Conference Series, 2012
http://iopscience.iop.org/1742-6596/396/4/042057

3. Rebootless Linux Kernel Patching with Ksplice Uptrack at BNL
Christopher Hollowell et al – Journal of Physics: Conference Series, 2012
http://iopscience.iop.org/1742-6596/396/4/042028