Jefferson Lab
High Performance and Scientific Computing

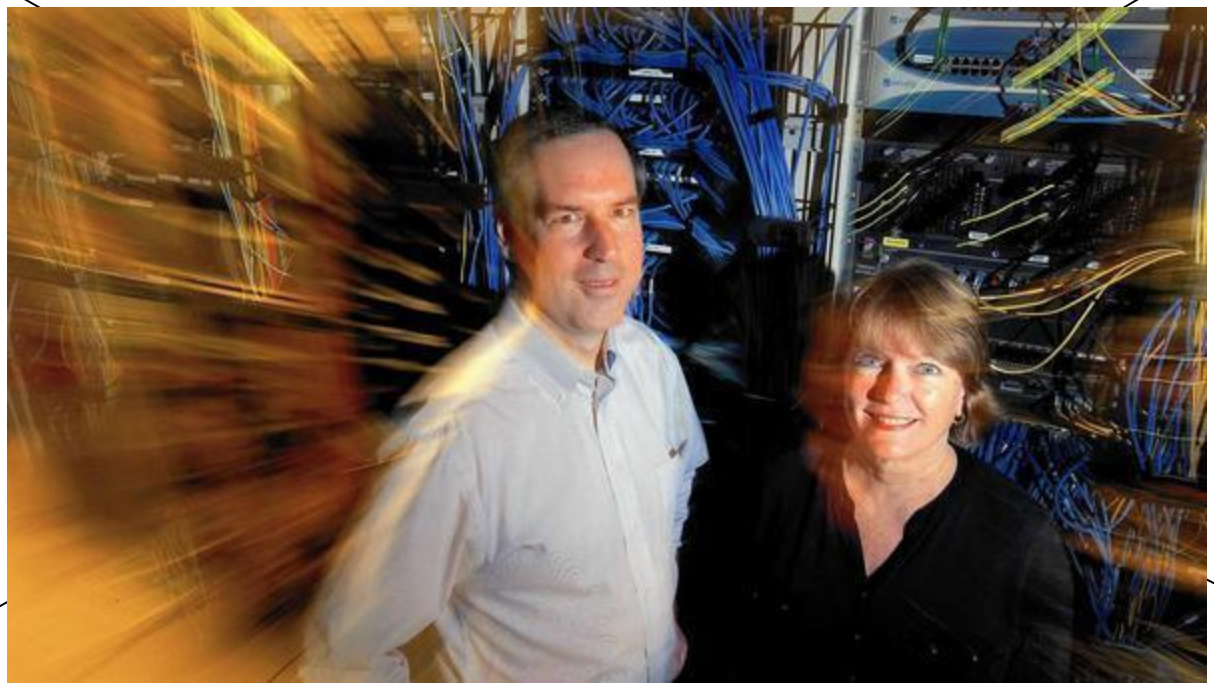

Sandy Philpott

HEPiX LBNL

October 17, 2016

Photo; Rob Ostermaier/Daily Press

Updates since DESY

- CEBAF 12GeV Physics support
- Computing
  - KNL, Broadwell installs
  - JLab in now an Intel Parallel Computing Center
  - Using Salt to build and manage systems
- Storage
  - Lustre 2.5.3 to IEEL 2.4.1 upgrade - June
  - 3 8TB*40 disk storage servers install underway
  - LNet routing - added IB->OPA
- Facilities
  - Data Center ongoing work continues …
- Looking ahead…

Scientific Computing

CEBAF 12GeV Physics support

- Added 2 more10gigE gateways from online to offline system
- Next computing and software readiness review Nov 10-11

We have begun using SALT in production to build and manage systems

Experimental Nuclear Physics hardware procurement:

64 Dual E5-2697V4 18 core 2.3GHz, 64GB RAM, FDR IB

Broadwell compute nodes for

- ExpPhy (48) compute farm, Hall D online (8), Hall C project (8)
- These nodes will run CentOS 7.2
  - Current farm nodes will be upgraded from CentOS 6.5 as the new nodes fill with jobs

HPC Computing

DOE LQCD-extII project: USQCD

3 sites: JLab, FNAL, BNL

FY16 procurement installed at Jlab; ~ $1M

Investigated several possibilities …

- Intel Xeon Phi / Knights Landing
  - Single socket, self hosting, largest on-package memory, >64 cores
- NVIDIA Pascal GPU, CUDA
- Intel Broadwell CPU server

Consideration factors

- hardware availability timeline
- high speed network – 100 Gbps price/performance; OmniPath or IB
- reflective benchmarks
- available configurations

See Wednesday's talk "Jlab's SciPhi-XVI" Knights Landing Cluster"

Lustre Disk Hardware

2016 hardware purchase – require 12 MB/s/TB

3 dual Xeon E5-2630V4 10 core 2.2GHz, 256GB RAM, Avago SAS 9300-8e 12Gb/s HBA; 40*8TB HGST; QDR IB

2015 Servers

2 dual Xeon E5-2630v2 6 core 2.6GHz, 128GB RAM, SAS3, 40*8TB Hitachi Ultrastar, 3*400GB Seagate SSD, LSI 9300-8e HBA, QDR on motherboard, FDR add-on

- RAID-Z, JBOD
- Fully redundant – 2 shelves connect 2 to hosts; currently installed non-HA
- SSDs (not yet in production use)

2014 Servers

4 dual Xeon E5-2630v2 CPUs, 30*4TB and 4*500GB SATA Enterprise disk drives, LSI 9361-8I RAID Controller with backup, 2*QDR ConnectX3 ports

- With RAID-Z, JBOD …no longer need hardware RAID
- Stable after ongoing issues with CATERRS
  - upgraded BIOS and disk controller firmware, set zfs_arc_max to default

Metadata servers in production since 2015

- 2 Dell R720s, E5-2620 v2 2.1GHz 6C, 64 GB RDIMM, 2 * 500GB 7.2K SATA
- PowerVault MD3200 6G SAS, dual 2G Cache Controller, 6 * 600GB 10K disk
  - ldiskfs

Lustre Filesystem

Intel Enterprise Edition Lustre 2.4.1 installed in June

- Upgraded from community Lustre 2.5.3
- Required OSS upgrades from CentOS 6.5 to 6.7
- Required larger than 100GB /boot partition!
- All clients upgraded as well
- Now 2 PB filesystem, filling from 60% to 80%

- Ongoing issues
  - Lustre lost files?
    - from incomplete lsf find by OSTs upon decommissioning or outage?
  - Small I/O hurts performance
    - Need to find and squash offenders! lltop, …
- LNet routers added, for connecting to Omni-Path

  Dual Xeon E5-2630V4 2.2GHz 10 core, 64GB RAM, QDR, OPA

Facilities Update

To meet DOE goal of PUE of 1.4, power and cooling are being refurbished in 2015-16 (announced last year, slow implementation…)

- New 800 KW UPS (done)
- 3 new 200 KW air handlers (+ refurbished 180)
- All file, interactive, infrasture servers will move to dual fed power, one side of which will be generator backed (99.99% uptime)

Transitions yet to come

- Rolling cluster outages to relocate and re-rack to 18-20 KW/rack as opposed to 10-12 KW today (starts in Nov/Dec)
- Anticipate 2 days outage per rack (3-4 racks at a time) plus 4 days full system outage over the next 8 months, so <2% for the year

Looking ahead

SSDs into production in the Lustre fileservers …

Computer Center Efficiency Upgrade and Consolidation continues…

CentOS 7.2 support for Experimental Physics computing

Installation of second tape library, as growth exceeds current library with 12GeV accelerator and experiments - ~2018 timeframe?

Data management, mining, indexing for Physics discovery …

Exascale Lattice Gauge Theory Opportunities and Requirements for Nuclear and High Energy Physics

https://www.eurekalert.org/pub_releases/2016-09/ddoe-tec090716.php