Site Report

Thomas Jefferson National Accelerator Facility
https://scicomp.jlab.org/docs

Sandy Philpott

HEPiX UCSD/SDSC

March 25, 2019

Highlights since Fall18 PIC

- CEBAF 12GeV beam to all 4 Experimental Halls
  - SSD gateway improves I/O from CEBAF online DAQ to MSS
  - Computing and software readiness review held November 2018
    - Positive outcome
- Computing
  - Slurm in full production
  - HOW2019 just wrapped up last week at Jlab
    - HSF, OSG, WLCG joing meeting, OSG tutorials
    - Looking forward to the next 10+ years of Scientific Computing
  - Hardware installations
    - GPUs for USQCD
    - Repurpose 7-year-old USQCD resources to JLab
      - Resource sharing update for Theory, Experimental Physics
  - Offsite resources accessible
- Disk Storage
  - Lustre 2.5 -> 2.10 upgrade underway
  - 4 Lustre OSS
- Tape Storage
  - New IBM TS4500 on order
  - lin_tape driver status
- And a Farewell !

Computing

A new 19g GPU cluster installation is underway

32 servers, each with

8 RTX 2080 GPUs

SWIF – JLab workflow tool

updating to support submissions to offsite resources

OSG, NERSC

Slurm is in full production for USQCD and ENP

from PBS/Torque/Maui

Offsite Computing Resources

- Open Science Grid
  - OSG submit host is updated from 1 gigE to 10 gigE
  - Planning for redundancy in services
  - GlueX (Hall D) in production
  - CLAS12 (Hall B) and others in consideration
  - HOW2019 participation key for JLab staff and Users
- NERSC
  - Our SWIF workflow tool eases submissions and Globus data transfers for users, and integrates with our local compute farm
    - Still, some problems with bottlenecks still are under investigation
- Cloud Services, still considered, as they are more cost effective and we need bursts …

Disks and Filesystems

Lustre and ZFS/NFS, 3+ petabytes

- 2019: replace Lustre MDS
  - Dell: 2 R540 with MD1420
- Lustre upgrade from 2.5 to 2.10
  - Switch to ZFS on MDS from ldiskfs
  - Limit problematic user I/O
- Plan to add 4 Lustre OSS to seed new 2.10 filesystem
- Hardware zoning, to provide redundancy for failover while separating SSD traffic from HDDs

Mass Storage

IBM TS3500 Tape Library in production

- 30 petabytes in 11 frames
- 8 LTO-8, 4 LTO-7, 8 LTO-6, 4 LTO-5 drives

IBM TS4500 library on order

- Additional 8 LTO-8 drives
- Relocate one S54 frame from current TS3500

IBM lin-tape driver issue

- Page allocation failures
- working with IBM to resolve

In Summary

SSDs into production for raw beam data to MSS improves I/O

Slurm environment in full production, for JLab resource sharing between Theory and Experimental Physics computing

Lustre 2.5 to 2.10 upgrade underway

Increase offsite resource usage

NERSC, OSG, and also at cloud and other providers

Installation of the new USQCD cluster: 32 8*RTX2080

Installation of IBM TS4500

HOW2019 meeting @ Jlab last week a huge success!

(a lead-in to holding CHEP ?)