Jefferson Lab LQCD Computing

April 2021 All Hands Meeting

Robert Edwards

Structure of NP/LQCD & JLab/LQCD programs

- Organizational structure is that all computing at JLab is under (new) Computational Sciences & Technology Division (CST) with its own Assoc. Dir.

- JLab Theory group leader is "Assoc. Dir. for Computations and Nuclear Theory"
- LQCD@JLab started in ~FY01 - all funds directed through JLab/Theory
  - Theory and CST group have worked closely over 20 years
  - CST division very responsive to needs of LQCD Theory

- Systems deployed based on computational requirements for LQCD Theory program & serve LQCD community as a whole
  - Program is coordinated with USQCD and Scientific Program Committee
  - Funding directly from NP

- Similar structure for Experimental Nuclear Physics at JLab & computing
  - Expt. Physics Div. receives funds with computing systems operated by CST

Nuclear and Particle Physics Computing Initiative

NPPLCI started in 2018 - PI for project was Chip Watson, now Robert Edwards

- JLab - single purpose lab, NP funded, serves all of USQCD, and is complementary to the modified (2 lab, IC based) HEP LQCD project
- $1M per year, about half hardware, half labor
  (equals average NP investment per year at JLab for last 10 years, so no real change in funding for the lab)
- FY2018: upgraded Jlab's KNL resources (added to system from FY2016)
- FY2019: upgraded Jlab's GPU resources (gamer-card system)
- FY2020/21: new AMD CPU + GPU system

More details in talks by Bryan Hess and Amitoj Singh

Nuclear and Particle Physics Computing Initiative

- Going forward, DOE asked JLab to merge LQCD into JLab Theory Group Annual reporting structure
  - First report (in Feb. 2020) included highlights of scientific program and future plans
  - Expect to carry this structure going forward

- DOE responds to the reports; however, not under an annual panel review process

- Instead, anticipate a broad review at the 4-year NP Theory Comparative reviews
  - Version in 2013 was across all of NP Theory (labs + universities)
  - 2016 review was only lab programs
  - Expect soon (2021?) at least a review of lab theory programs

JLab & software development

- Heavy involvement since 2001 & SciDAC-1, then SciDAC-2 & 3
- Currently, lead institution for ASCR/NP SciDAC-4 project
  - involving MIT, BNL, ORNL, W&M, GWU, MSU, UNC, LANL + industry
- Also direct NP portion under Exascale Computing Project

- Programs allow LQCD to leverage our local/commodity resources into efficiently utilizing national resources
  - This is how LQCD & NP in particular got onto ORNL NVIDIA systems early on
  - But these days, GPU-s are the main game in town

- Frontier+Aurora have pushed the “s” in GPU-“s”
  - Opportunity/necessity to diversify our codes
- Opportunity for LQCD to exploit for new commodity-based systems

Exascale Computing Project

- At some point, became clear
  - [Not Exascale] Perlmutter - AMD CPU + NVIDIA GPU + Cray network - now
  - Frontier - AMD CPU + AMD GPU + Cray network - soon
  - Aurora - Intel CPU + Intel GPU + Cray network - not as soon

- Thrusts for code bases
  - NVIDIA widely supported - new activities devoted to “refactorizing”
  - QUDA refactorized - now “backend” support for NVIDIA+AMD+Intel GPUs
  - Grid support for all, but focused more on Intel
  - CPS & MILC - rely on QUDA, but more focused towards Intel
  - NDA-s huge restriction - limit informed public discussions

- NP & Chroma - refactored to support more GPU systems, but prioritized AMD
  - Support now for AMD+NVIDIA quite mature - some parts ***production ready***
  - Software available on GitHub

New LQCD system acquisitions

- JLab carried out an “Alternatives Analysis” and finished in early 2020
  - Gamer GPU system served limited set of applications
  - Input from community - need to expand portfolio of applications
  - AMD GPU an emerging technology in addition to NVIDIA
  - Consider two main metrics to quantify price-performance
    - Wilson dslash - proxy for inverters
    - ZGEMM based contractions considering possible swapping to main memory

- Development under ECP & SciDAC allowed for benchmarks across different systems
- Selected an AMD based system

See more details in talk by Amitoj Singh

LQCD @ JLab

https://lqcd.jlab.org - all project details, including usages here

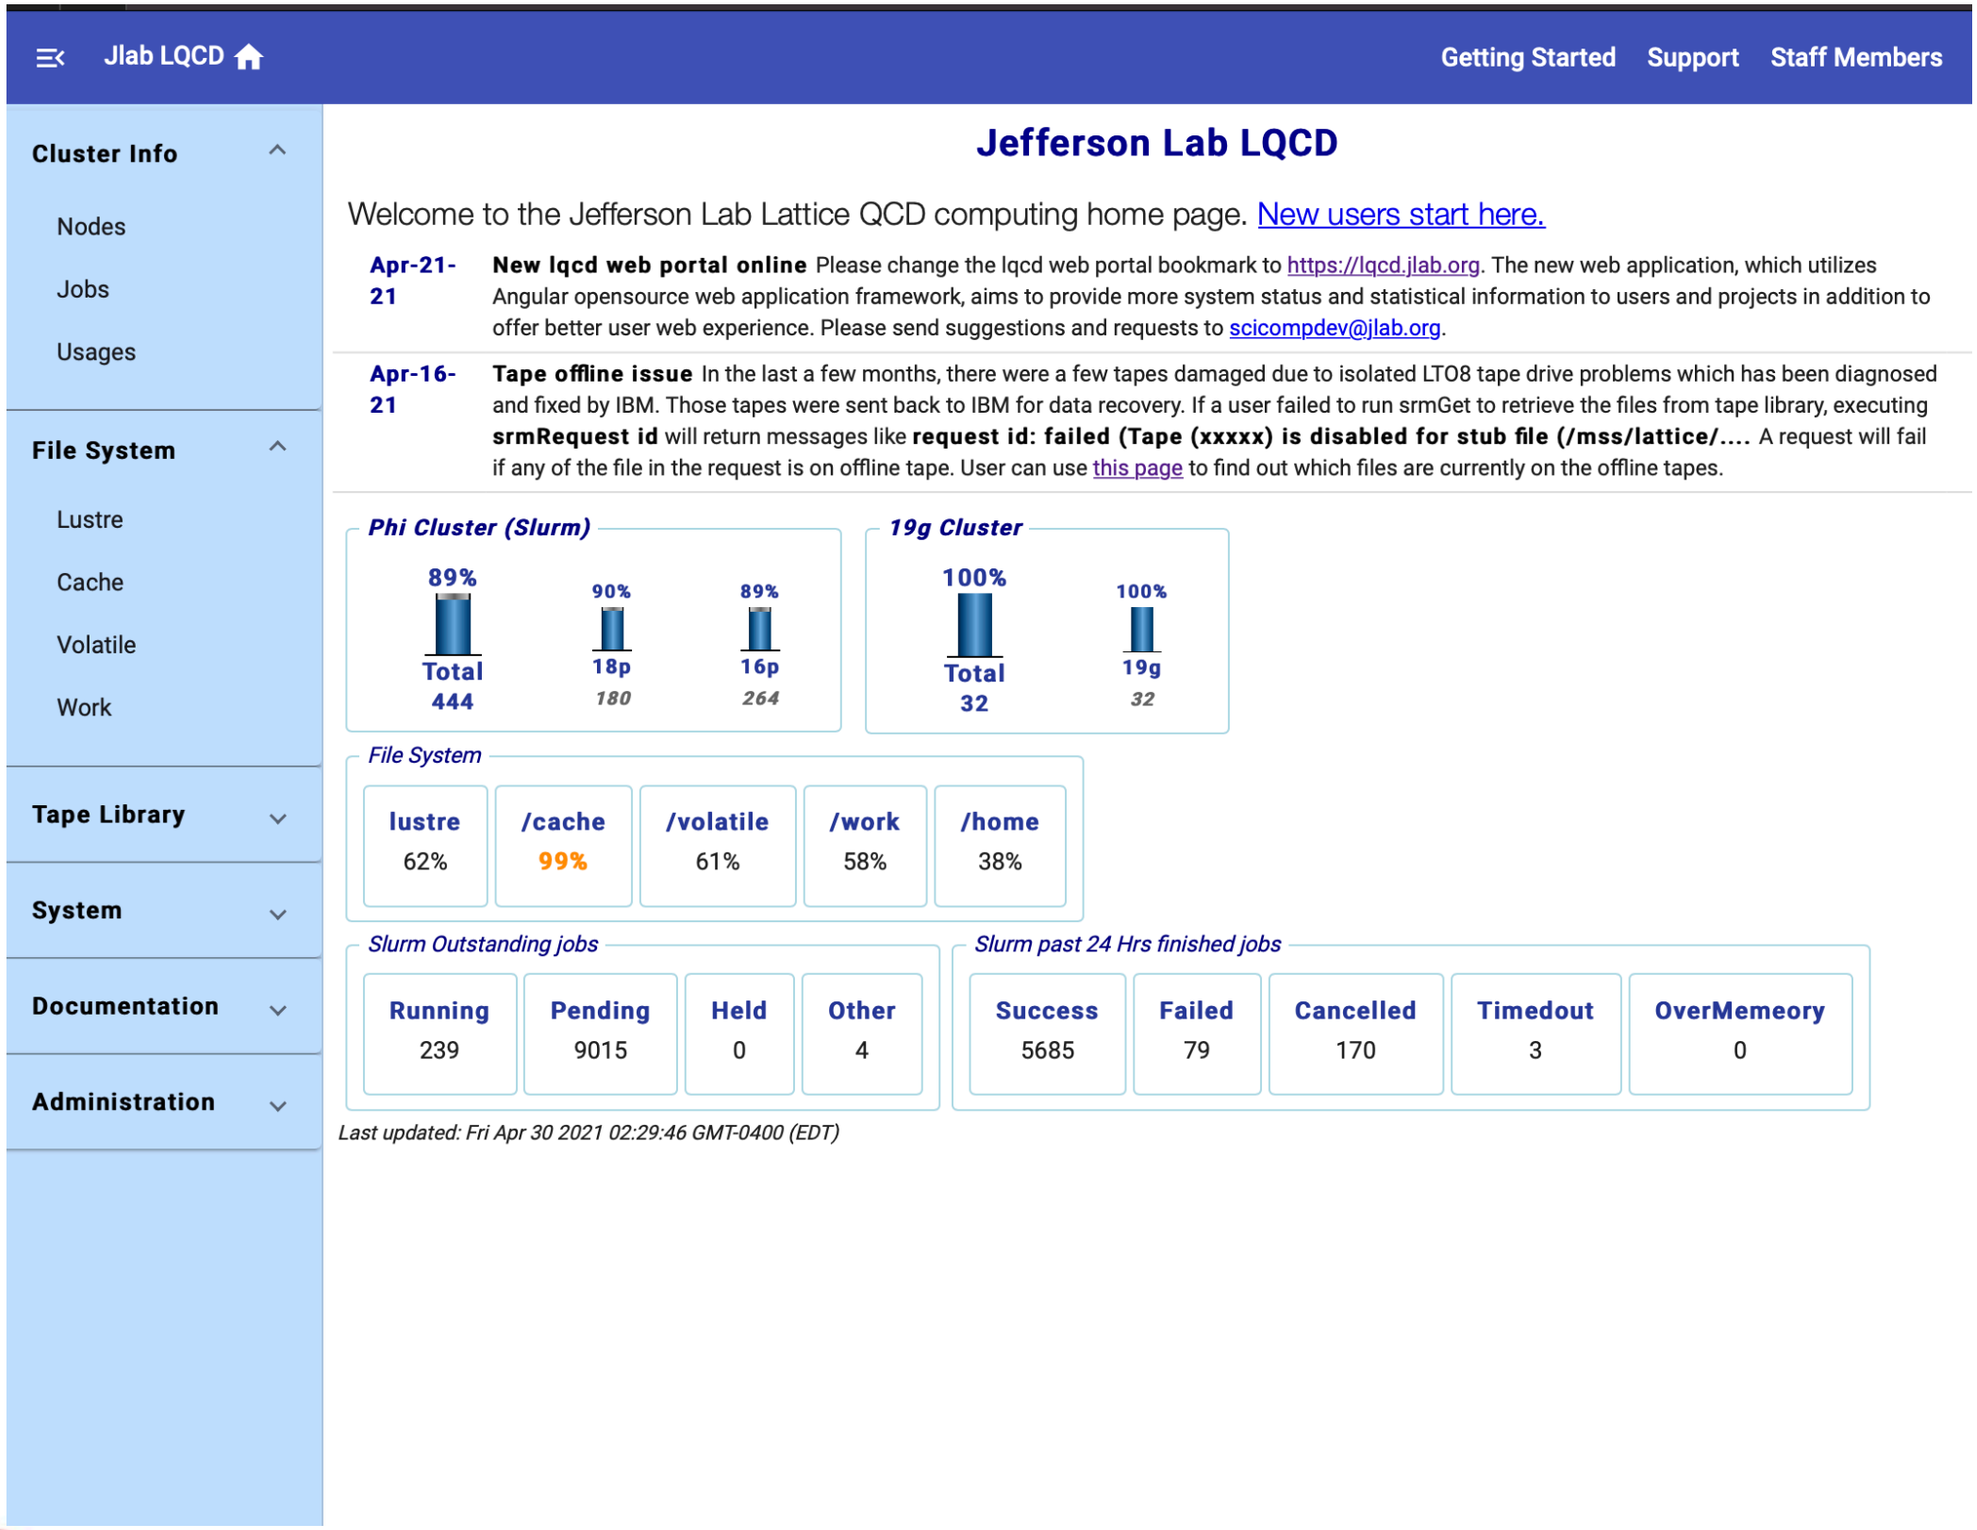

Next up is JLab LQCD operations

- Organizational structure:
  - Computational Sciences & Technology
    - Graham Heyes - Head of Scientific Computing - overseas operations of LQCD
    - Amitoj Singh - LQCD Site Manager - in April assumed position from Bryan Hess
    - Jie Chen - Performant Cluster Scientific Computing Group Leader (PCSCI)
    - Frank Winter - PCSCI
    - Eloy Romero - PCSCI
  - Theory
    - Robert Edwards - NP initiative manager & Deputy Theory Group Leader

- Presentations by Bryan Hess (operations) & Amitoj Singh (new purchases)