LQCD Site Report

Jefferson Lab

Amitoj Singh

Thursday, April 21, 2022

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

Topics covered in this talk

- JLab Organization and your support team
- User documentation & support
- State of compute and storage
- A few operations highlights
- Allocations summary and SLURM job trends
- Some tips on using the cluster resources effectively
- Future hardware acquisition
- Questions

Jefferson Lab Organization – hiring and growing


Amber Boehnlein

Computational Sciences and Technology Ad/CIO
Amber Boehnlein

Amitoj Singh
LQCD Site Manager
HPC Specialist

Scientific Computing Department (SciComp)
Graham Heyes

Management and Information Services Department (MIS)
Kari Heffner

Computing and Network Infrastructure Department (CNI)
Andy Kowalski

Data Science Department Head
Malachi Schram

Experimental Physics Scientific Computing Infrastructure Group (EPSCI)
D Lawrence

Performant Cluster Scientific Computing Infrastructure Group (PCSCI)
J Chen

Information Resources (IR)
D Cochran

Scientific Computing Operations
B Hess

Cyber Security
G Norwicki

Networking/Telecom
B Morris

Systems
K Edwards

Helpdesk
J Perry


Amitoj Singh


Graham Heyes


Jie Chen

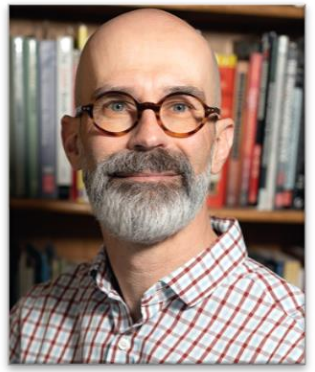
Bryan Hess


Andy Kowalski

Your JLab LQCD Support Team

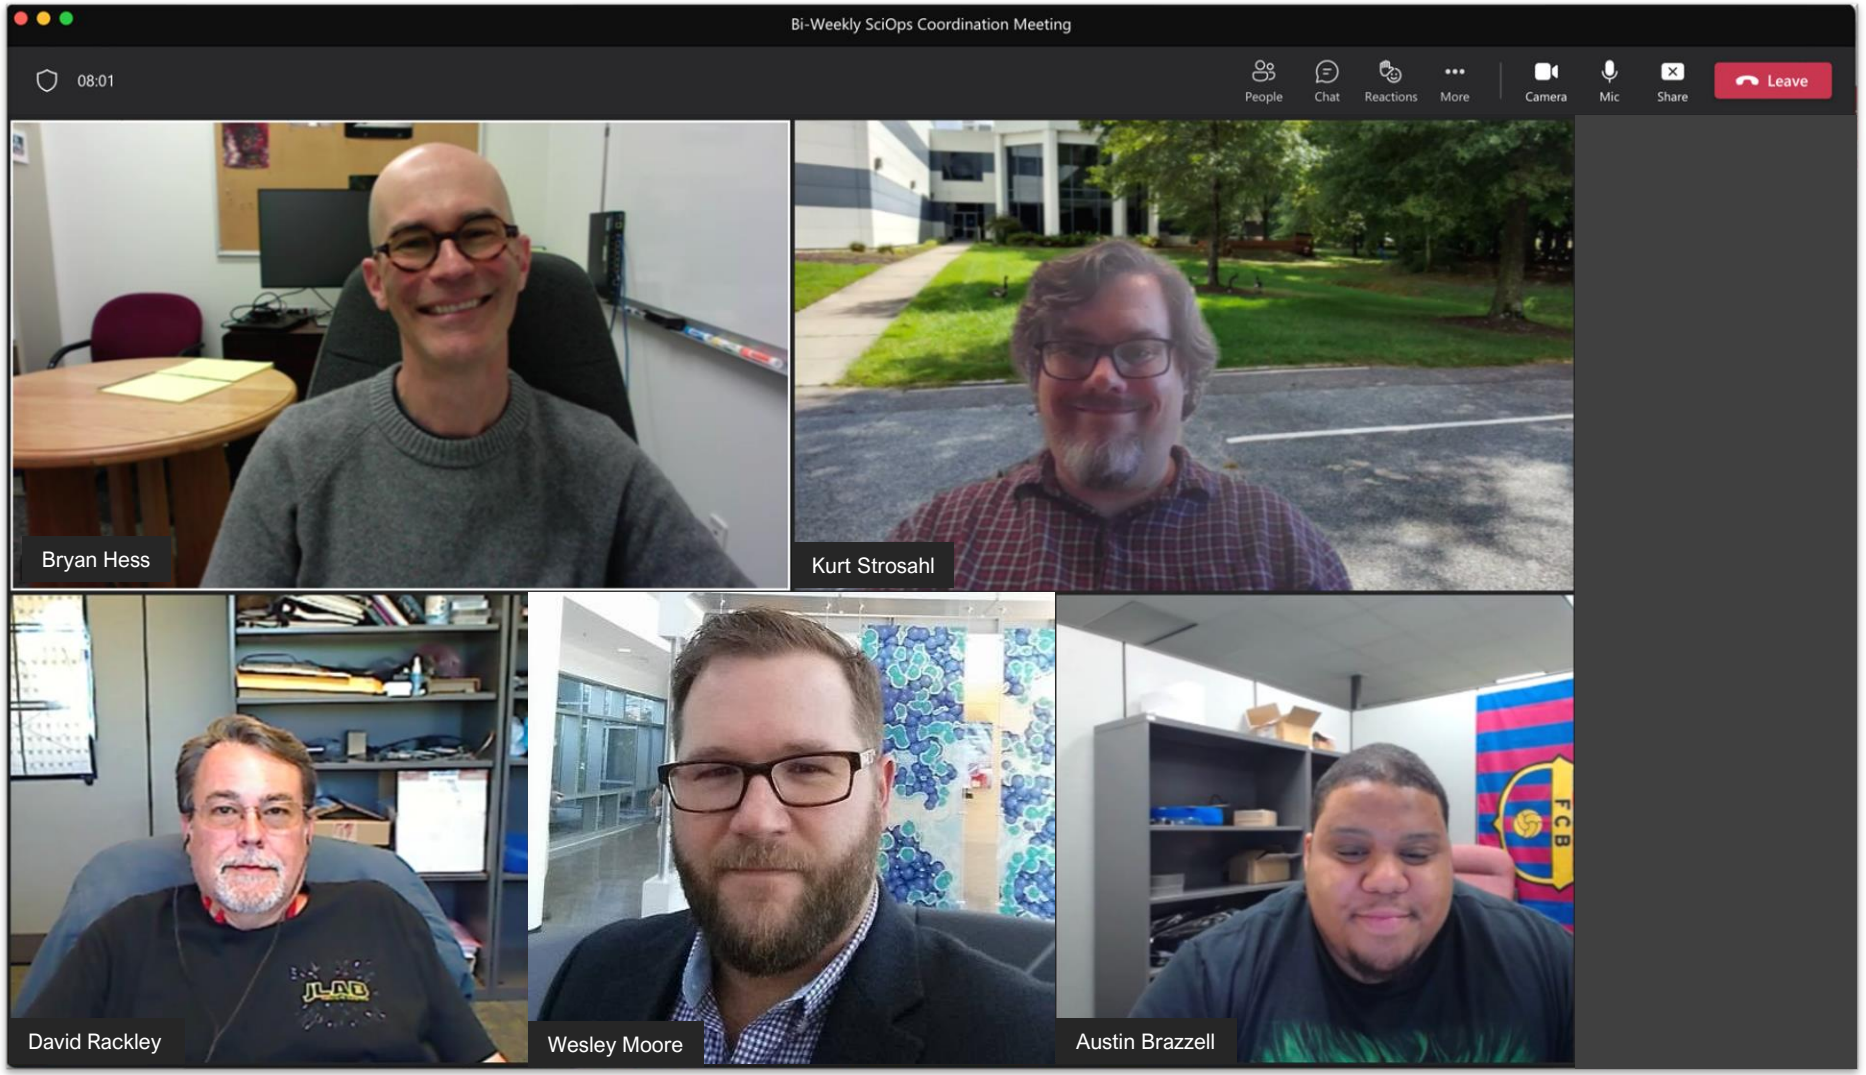

User Documentation & how to ask for support

https://lqcd.jlab.org

JLab Cluster Layout

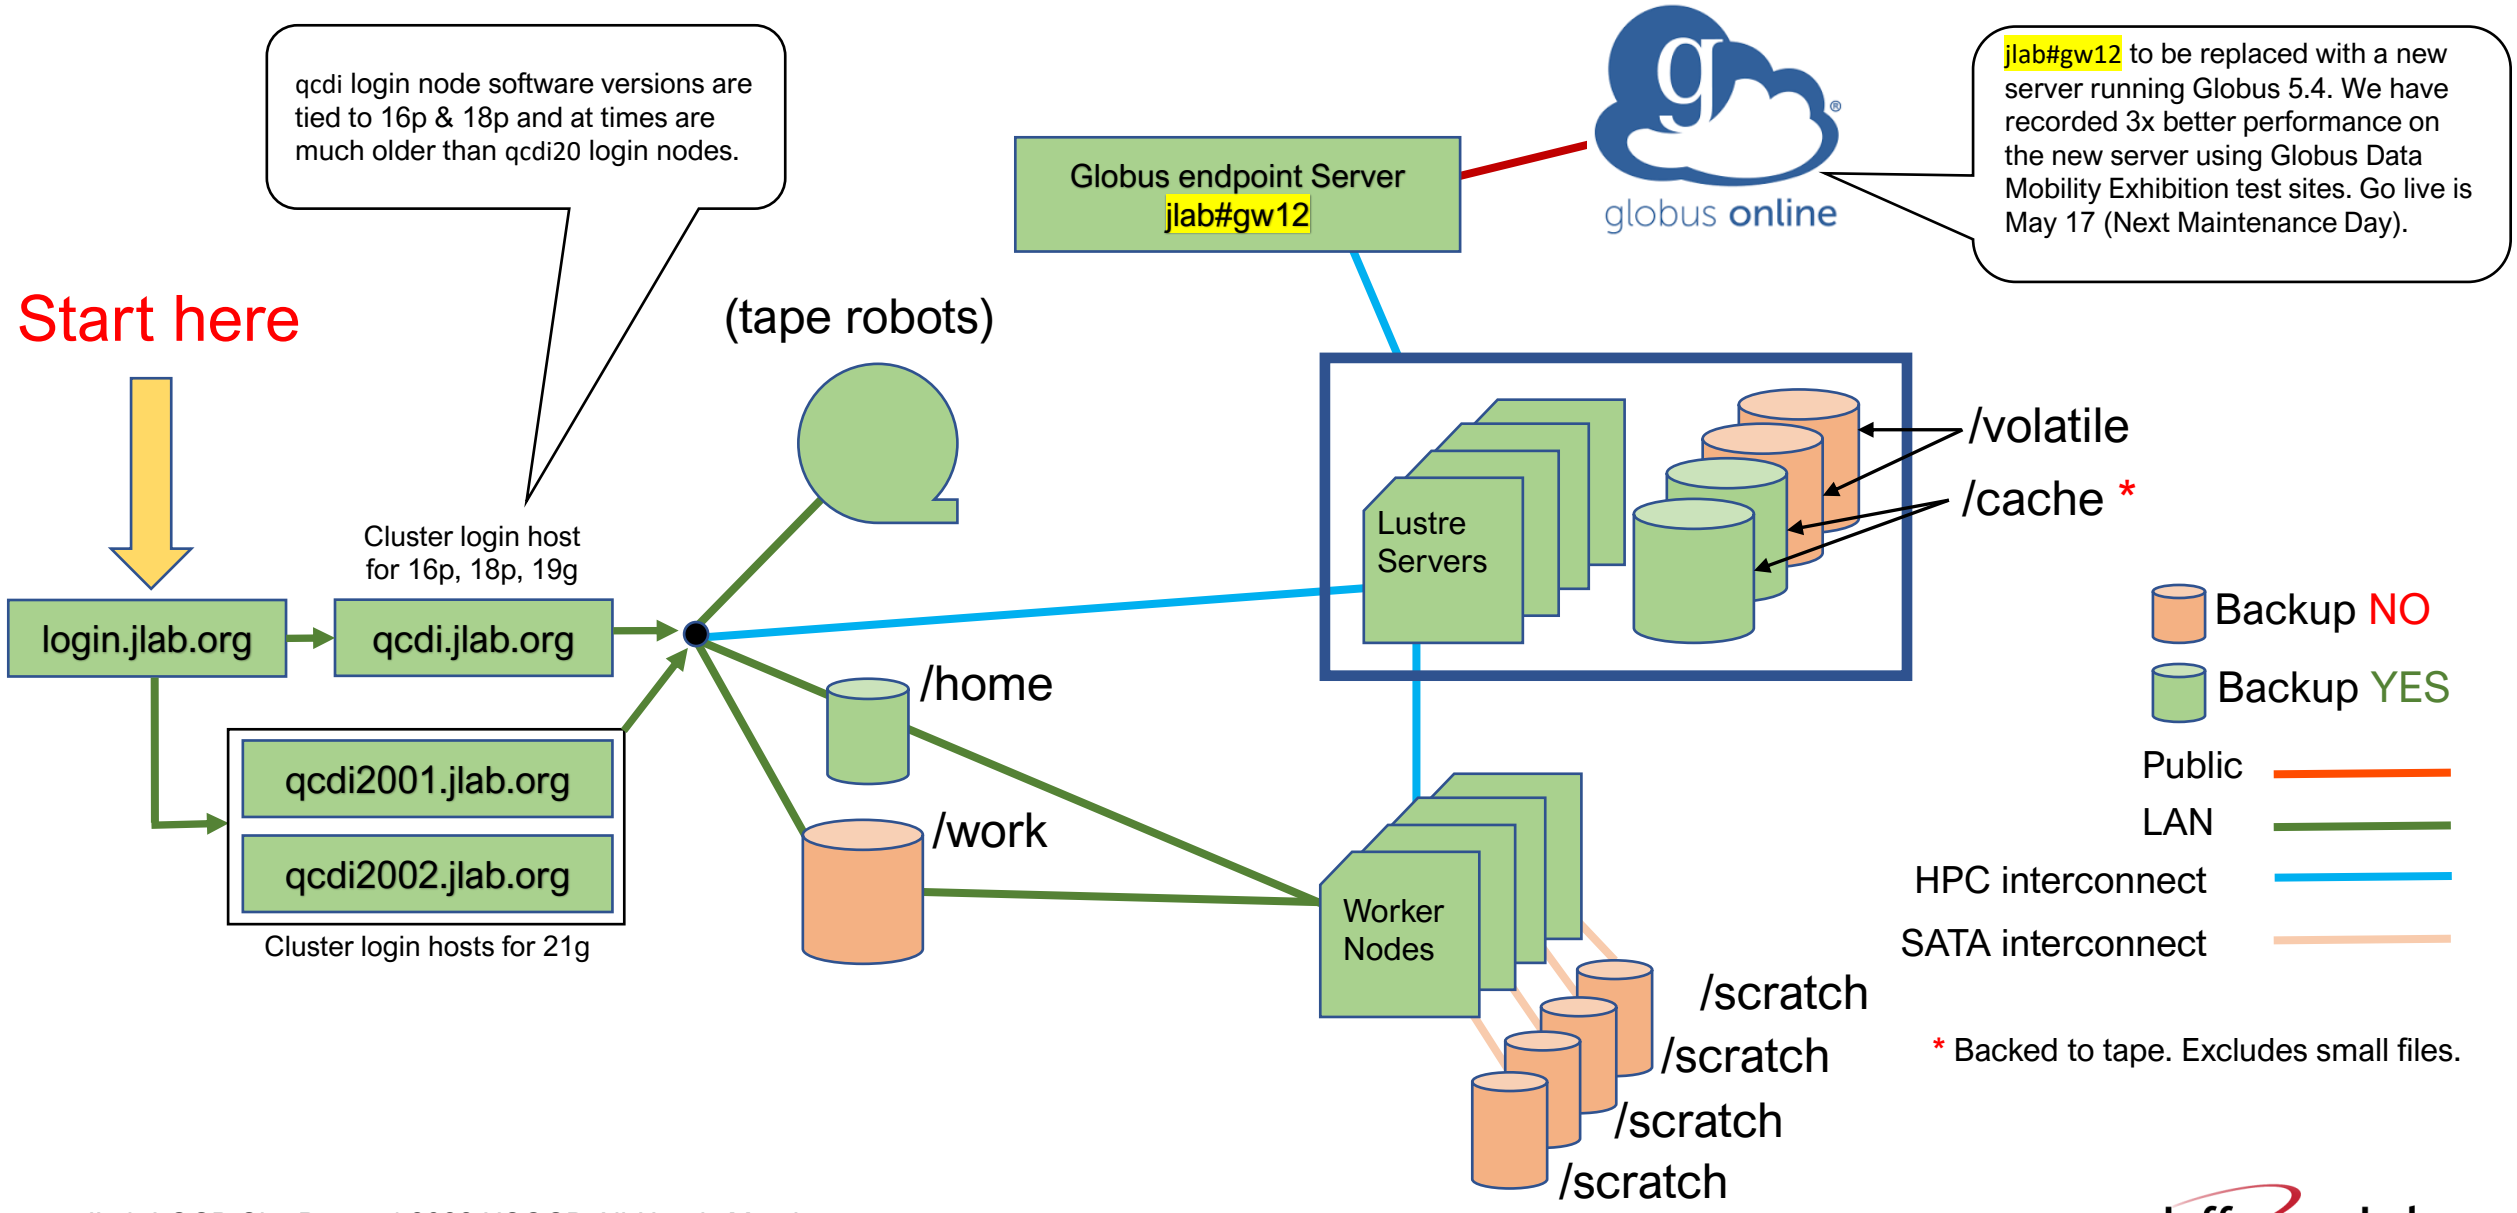

Current computing as pledged for 2022-23 USQCD allocations

- Compute
  - 432 node Xeon Phi / KNL cluster (“16p/18p”)
    - Single socket 64 core KNL (with AVX-512 8 double / 16 single precision) 192 (98) GB main memory / node 16p (18p)
    - 32GB high bandwidth on package memory (6x higher bandwidth)
    - 100 Gbps bi-directional OmniPath network fabric (total 25GB/s/node) 32 nodes / switch, 16 up-links to core / switch
    - Total: 3.1 M node-hours = 198.4 M KNL-core-hours
  - 32-node GeForce GPU cluster (“19g”)
    - Eight-GPU RTX-2080 nodes
    - 8GB memory per GPU, 192GB memory per node. Each on 100g OmniPath Fabric
    - Total: 229.86 k node-hours = 1.84 M RTX2080-GPU-hours
  - 8-node AMD GPU Cluster (“21g”)
    - Eight-CPU AMD MI100 nodes with Inter-GPU Infinity interconnect. 32GB memory per GPU, 1TB memory per node.
    - Each on 100g InfiniBand Fabric
    - Total: 57.47 k node-hours = 0.46 M MI100-GPU-hours

Current disk and tape status

- Lustre - Storage for /volatile and /cache
  - 2.3PB (actual available 1.9PB) parallel and distributed Lustre file-system.
- NFS file server - /work and /home on ZFS.
  - In the process of configuring a new file server. Migration of existing data will commence soon followed by deployment.
  - /work is not backed up. Quotas are assigned per project.
  - /home is backed up weekly.
- Tape Storage
  - Majority data is on LTO M8 (9TB) media.
  - Since a year ago all writes are now to LTO8 (12TB) tape media.
  - **To date LQCD accumulated storage is 11PB on tape** (https://lqcd.jlab.org/lqcd/cacheDisk/project)
    - 9.7PB on lattice-p "permanent"
    - 1.3PB on lattice-t "temporary"
    - Tape storage for lattice-t USQCD (non-JLab) allocations are retained at Jefferson lab for 18 months after the allocation year ends, then the tapes are re-used.
  - All tape related costs (minus media) are covered by JLab.

A few operational highlights – 2021-22

- September
  - SLURM upgraded to version 20:
    - adds AMD GPU support for 21g + security fix.
    - LQCD distinct security realm for SLURM (was shared with farm).
    - Upgraded “on the fly” to avoid any outage or down time!
  - Legacy /home performance issue fix from 21g.
  - Two new interactive nodes established for 21g cluster:
    - qcdi2001 - MI50 on intel
    - qcdi2002 - MI50 on AMD
- October
  - 21g allocated and in production.
- November
  - Now monitoring individual GPU performance and temperature on all GPU clusters.
  - Consolidation of dead nodes and clean up of KNL clusters.
  - One of the 19g cluster worker node integrated ethernet NIC fail - added extra NIC as workaround since we have no extra systems. Motherboard ordered from KOI under warranty.
- December
  - 21g ROCm upgrade successful.
  - Fix for SLURM epilog/prolog race conditions that could kill jobs, other SLURM cleanup and puppet config.
- January
  - Replaced 19g cluster worker node motherboard after failure (KOI, on site).
  - 21g cluster worker node motherboard failure again, awaiting delivery. Atipa scheduled on site first week of Feb.
- March
  - 21g final node in service after significant hardware debugging and replacement.
- April
  - Multiple 21g ROCM upgrades as requested by users for performance and feature improvements.

2021-2022 Allocations Summary

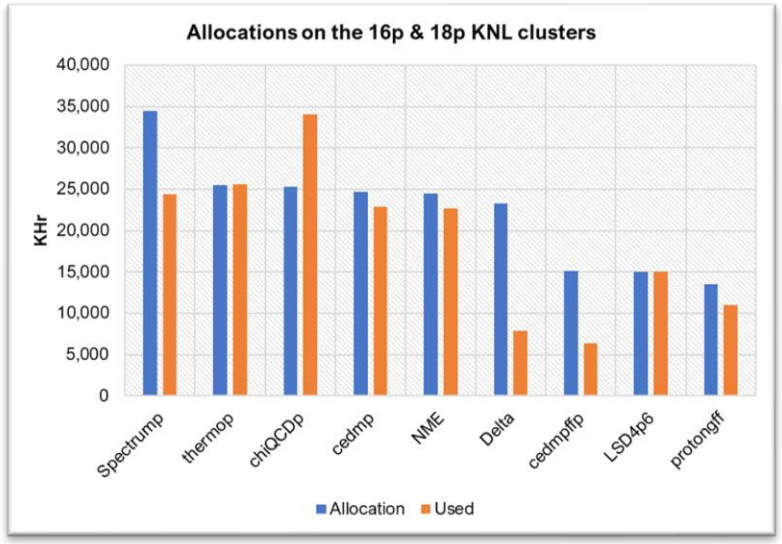

Total allocations used thus far with 83% of allocation year completed

85%

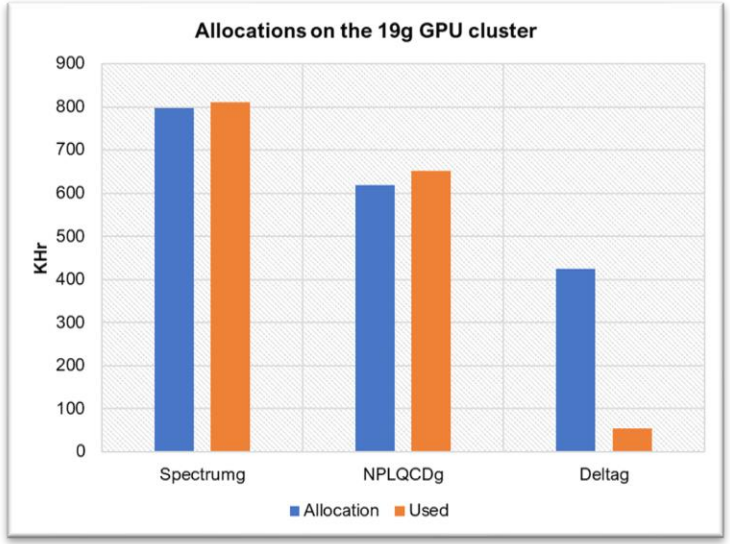

82%

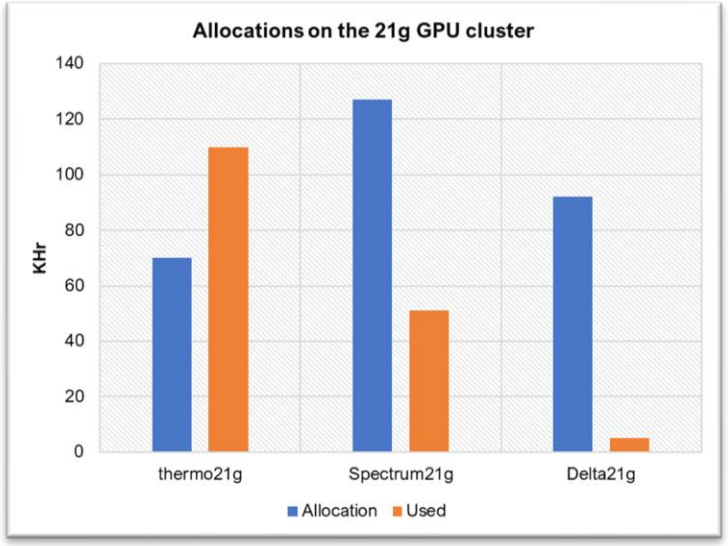

Total allocations used thus far with 78%* of allocation year completed

57%

KNL and NVIDIA GPU cluster allocations are on track. 21g cluster usage has been low due to the low number of allocated projects and a start date of Oct 1st vs Jul 1st. We have setup a `21g_dev` queue for quick turnaround on 21g for code development and optimization jobs esp. to encourage new users.

* Allocation start date of Oct-1-2021

PYTD (Jul-1-2021 till date) SLURM Job trends: Job Node Sizes

16p 18p KNL clusters

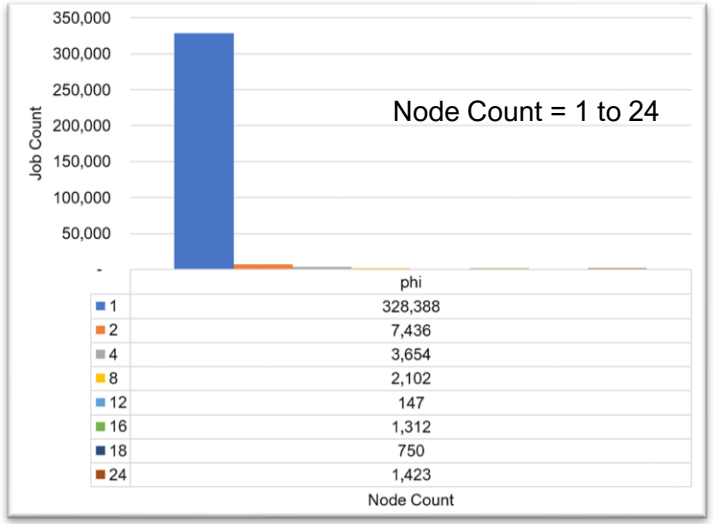

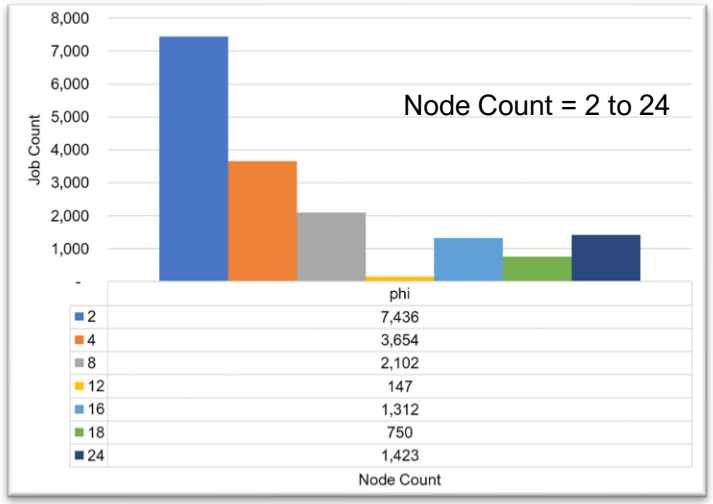

Single node jobs dominate 16p and 18p.

19g NVIDIA GPU cluster

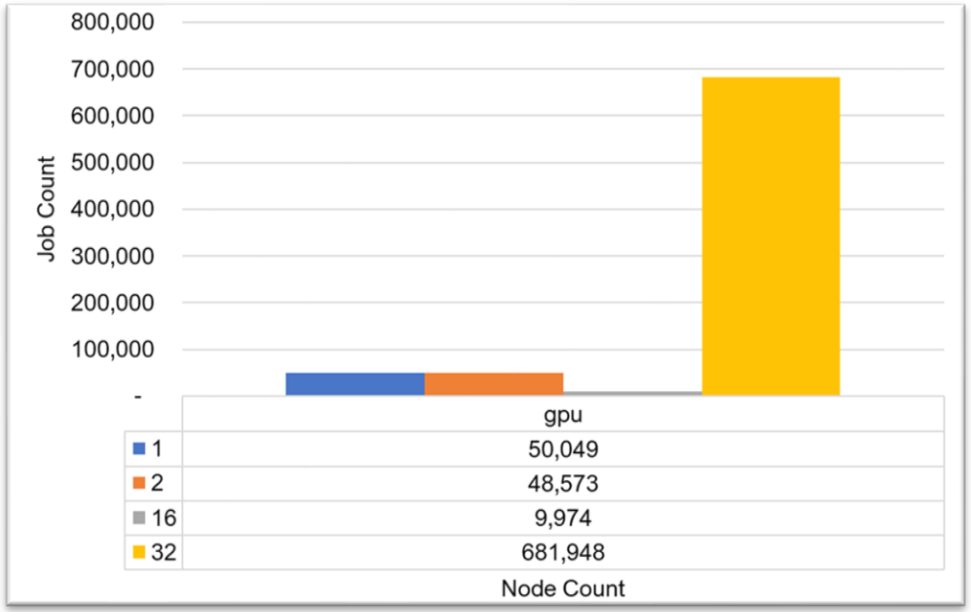

32 node jobs dominate 19g.

21g AMD GPU cluster

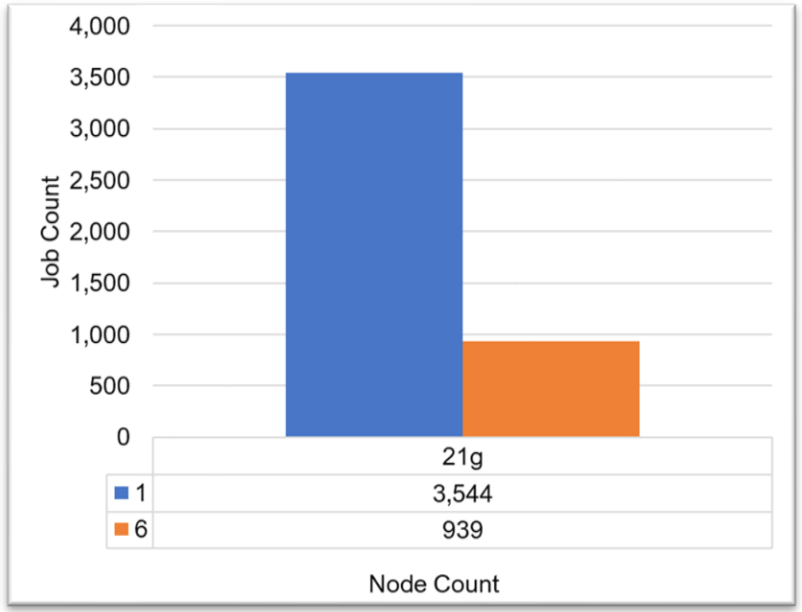

Single node jobs dominate 21g.

A few tips on compiler options

- In order to compile your code using CUDA, on the login nodes (qcdi1401 or qcdi1402) or 19g worker nodes check available CUDA versions as follows.

  $> module use /dist/modulefiles

  $> module ..

- A simple way to compile a kernel for MI100 on 21g: make sure to use

  --amdgpu-target=gfx906,gfx908

  which is like cuda_sm. gfx908 is for MI100, gfx906 is for MI50.

- More useful tips for 21g are here: https://scicomp.jlab.org/docs/using21g

```
[amitoj@qcdi1402 ~]$ module use /dist/modulefiles
[amitoj@qcdi1402 ~]$ module av

------------------------------------------------ /dist/modulefiles -------------------------------------------------
anaconda2/4.4.0   boost/1.78.0      cuda/9.0          gcc/7.2.0         go/1.15.4         singularity/3.7.4
anaconda2/5.2.0   cmake/3.17.5      curl/7.59         gcc/7.5.0         python3/3.8.7     singularity/3.8.3
anaconda3/4.4.0   cmake/3.18.4      gcc/10.2.0        gcc/8.4.0         python3/3.9.7
anaconda3/5.2.0   cmake/3.21.1      gcc/5.3.0         gcc/9.3.0         singularity/2.3.1
boost/1.74.0      cuda/10.0         gcc/7.1.0         go/1.13.5         singularity/3.6.4

---------------------------------------- /usr/share/Modules/modulefiles ----------------------------------------
dot         module-git  module-info modules     null        use.own

------------------------------------------------ /etc/modulefiles -------------------------------------------------
anaconda           ansys2020r1        gcc_4.8.2          gcc_5.3.0          gnuplot-5.2.4      mvapich2-1.8
anaconda2          ansys-old          gcc-4.8.2          gcc-5.3.0          gsl-1.15           mvapich2-2.1
anaconda3          cmake-3.13.4       gcc_4.9.2          gcc-6.2.0          hdf5-1.8.12        mvapich2-2.3a
anaconda-latest    curl-59            gcc-4.9.2          gcc-6.3.0          java_1.7           openmpi-1.10.2-hfi
ansys              gcc_4.6.3          gcc_5.2.0          gcc-7.1.0          java_1.8           zeromq
ansys18            gcc-4.6.3          gcc-5.2.0          gdb_7.11.1         mpi/openmpi-x86_64
[amitoj@qcdi1402 ~]$ module load cuda/10.0
[amitoj@qcdi1402 ~]$
[amitoj@qcdi1402 ~]$ nvcc --version
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2018 NVIDIA Corporation
Built on Sat_Aug_25_21:08:01_CDT_2018
Cuda compilation tools, release 10.0, V10.0.130
[amitoj@qcdi1402 ~]$
```

```
[amitoj@qcdi2001 ~]$ module av
---------------------------------------- /usr/share/Modules/modulefiles ----------------------------------------
dot  module-git  module-info  modules  null  rocmmod4.3.0  rocmmod4.5.2  rocmmod5.0.2  rocmmod5.1.1  use.own

------------------------------------------------ /etc/modulefiles -------------------------------------------------
mpi/mvapich2-x86_64  rocm/4.2.0  rocm/4.3.0  rocm/4.5.2  rocm/5.0.2  rocm/5.1.1

-------------------------------------------- /usr/share/modulefiles ---------------------------------------------
mpi/openmpi-x86_64  pmi/pmix-x86_64
[amitoj@qcdi2001 ~]$
[amitoj@qcdi2001 ~]$ module load rocm/5.1.1
[amitoj@qcdi2001 ~]$
[amitoj@qcdi2001 ~]$ hipcc --amdgpu-target=gfx906,gfx908 -o HipccHellWorld HipccHelloWorld.cpp
[amitoj@qcdi2001 ~]$
[amitoj@qcdi2001 ~]$ ./HipccHelloWorld
 System minor 0
 System major 9
 agent prop name
input string:
GdkknVnqkc

output string:
HelloWorld
Passed!
[amitoj@qcdi2001 ~]$
```

Future Hardware Acquisition

- None is planned for current FY22 and will reevaluate in FY23.
- BNL and FNAL machines coming online in CY2022 and CY2023.
  - Sufficient GPU resources in the portfolio of hardware available at all 3 sites.
- CY2024 promises exciting new hardware options such as:
  - CPU GPU Hybrid
    - AMD
    - NVIDIA
  - ARM
  - RISC-V
- By carrying hardware acquisition funds over multiple FY boundaries we can purchase a larger machine that best serves the needs of the USQCD collaboration.

Amitoj Singh

amitoj@jlab.org

Questions?

- Topics covered in today's talk
  - JLab Organization and your support team
  - User documentation & support
  - State of compute and storage
  - A few operations highlights
  - Allocations summary and SLURM job trends
  - Some tips on using the cluster resources effectively
  - Future hardware acquisition