

Lattice QCD Site Report

Jefferson Lab

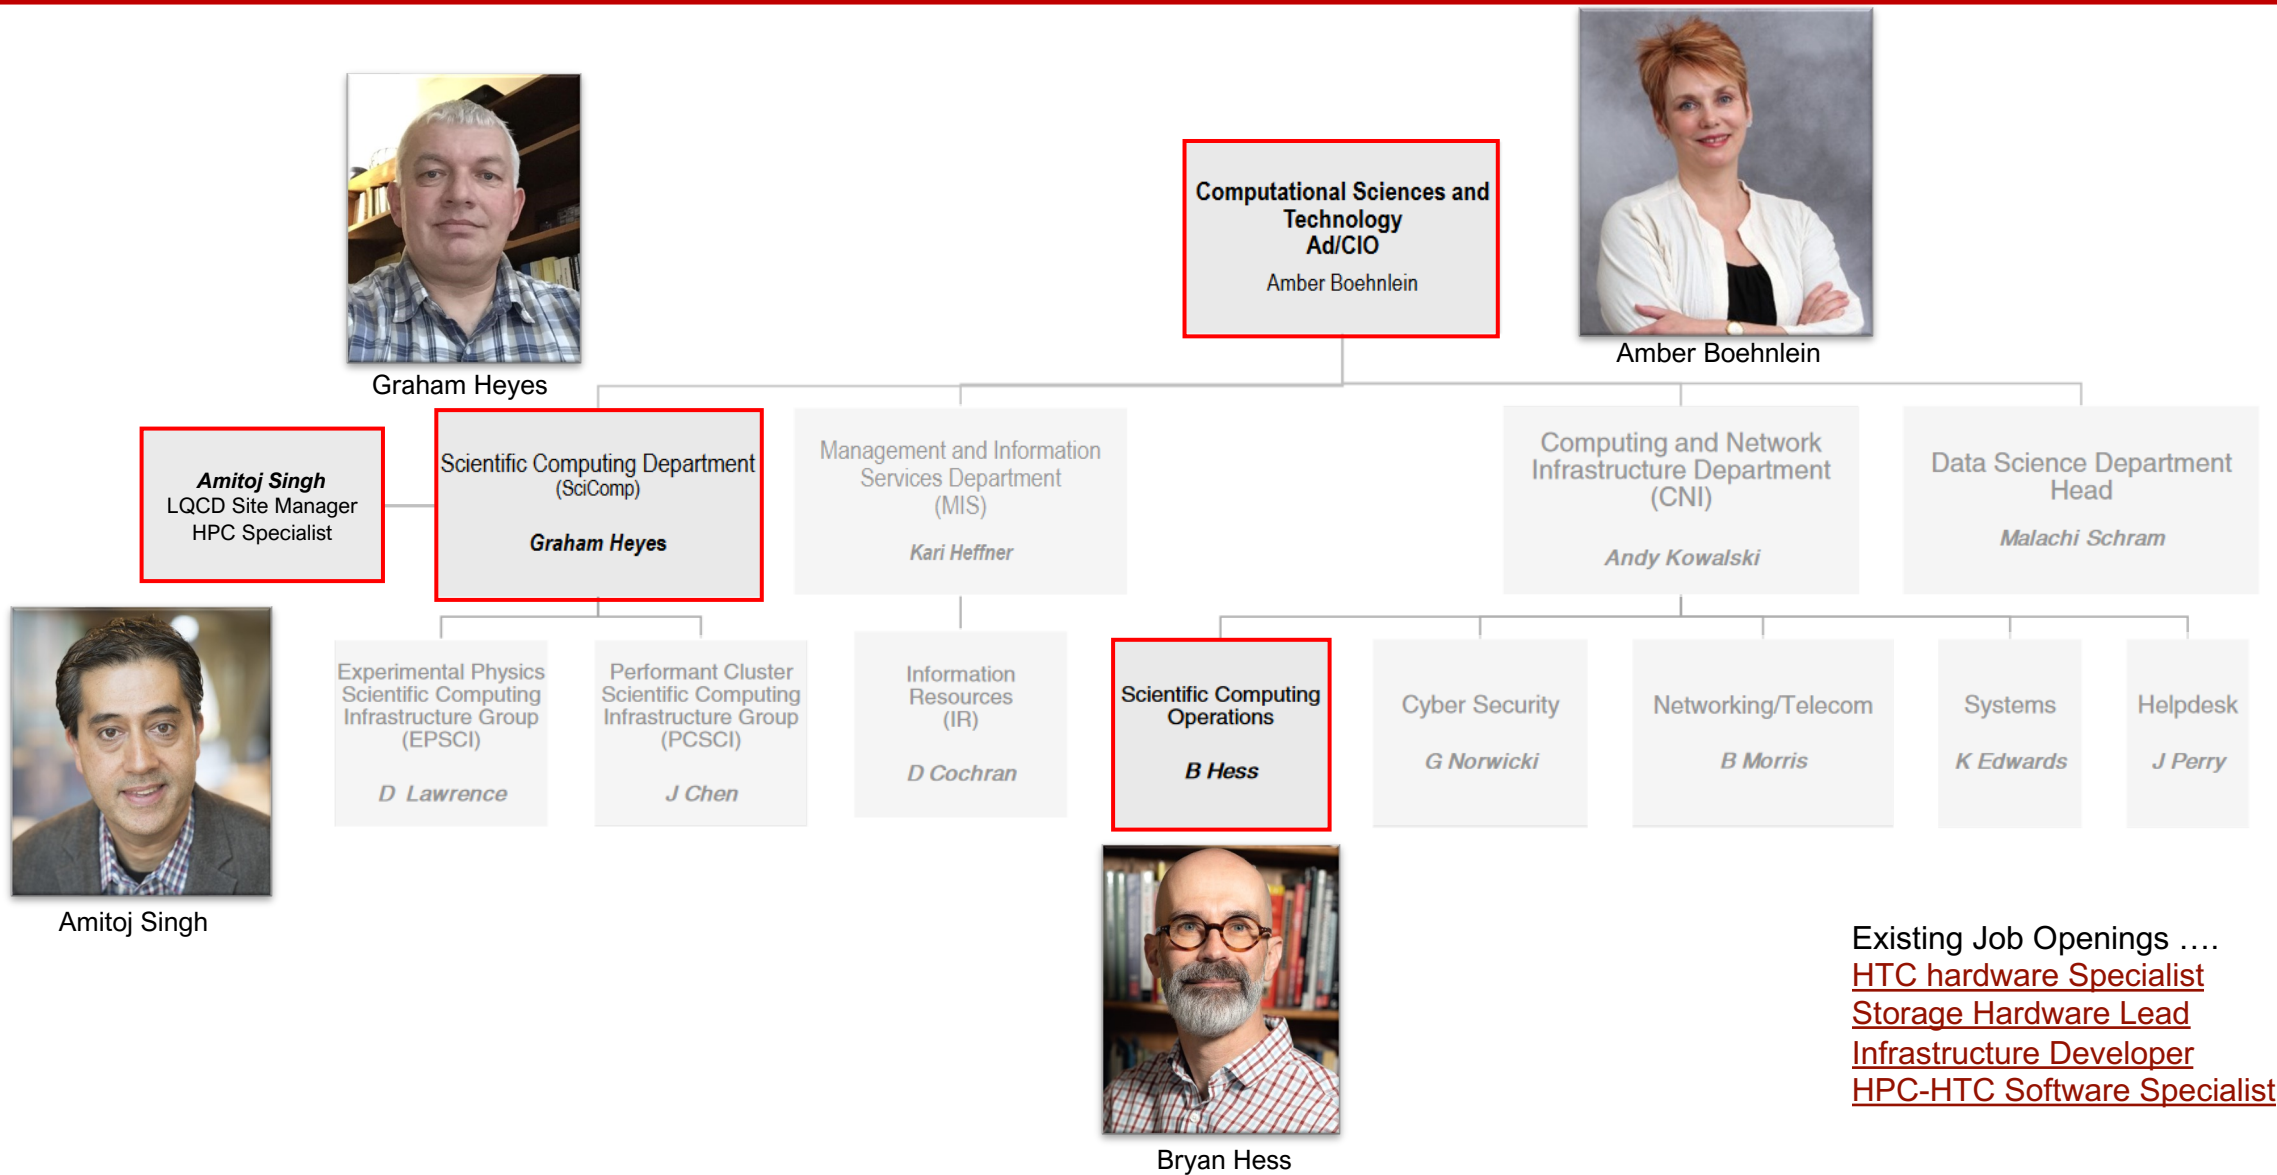
Amitoj Singh

Thursday, April 20, 2023

Jefferson Lab

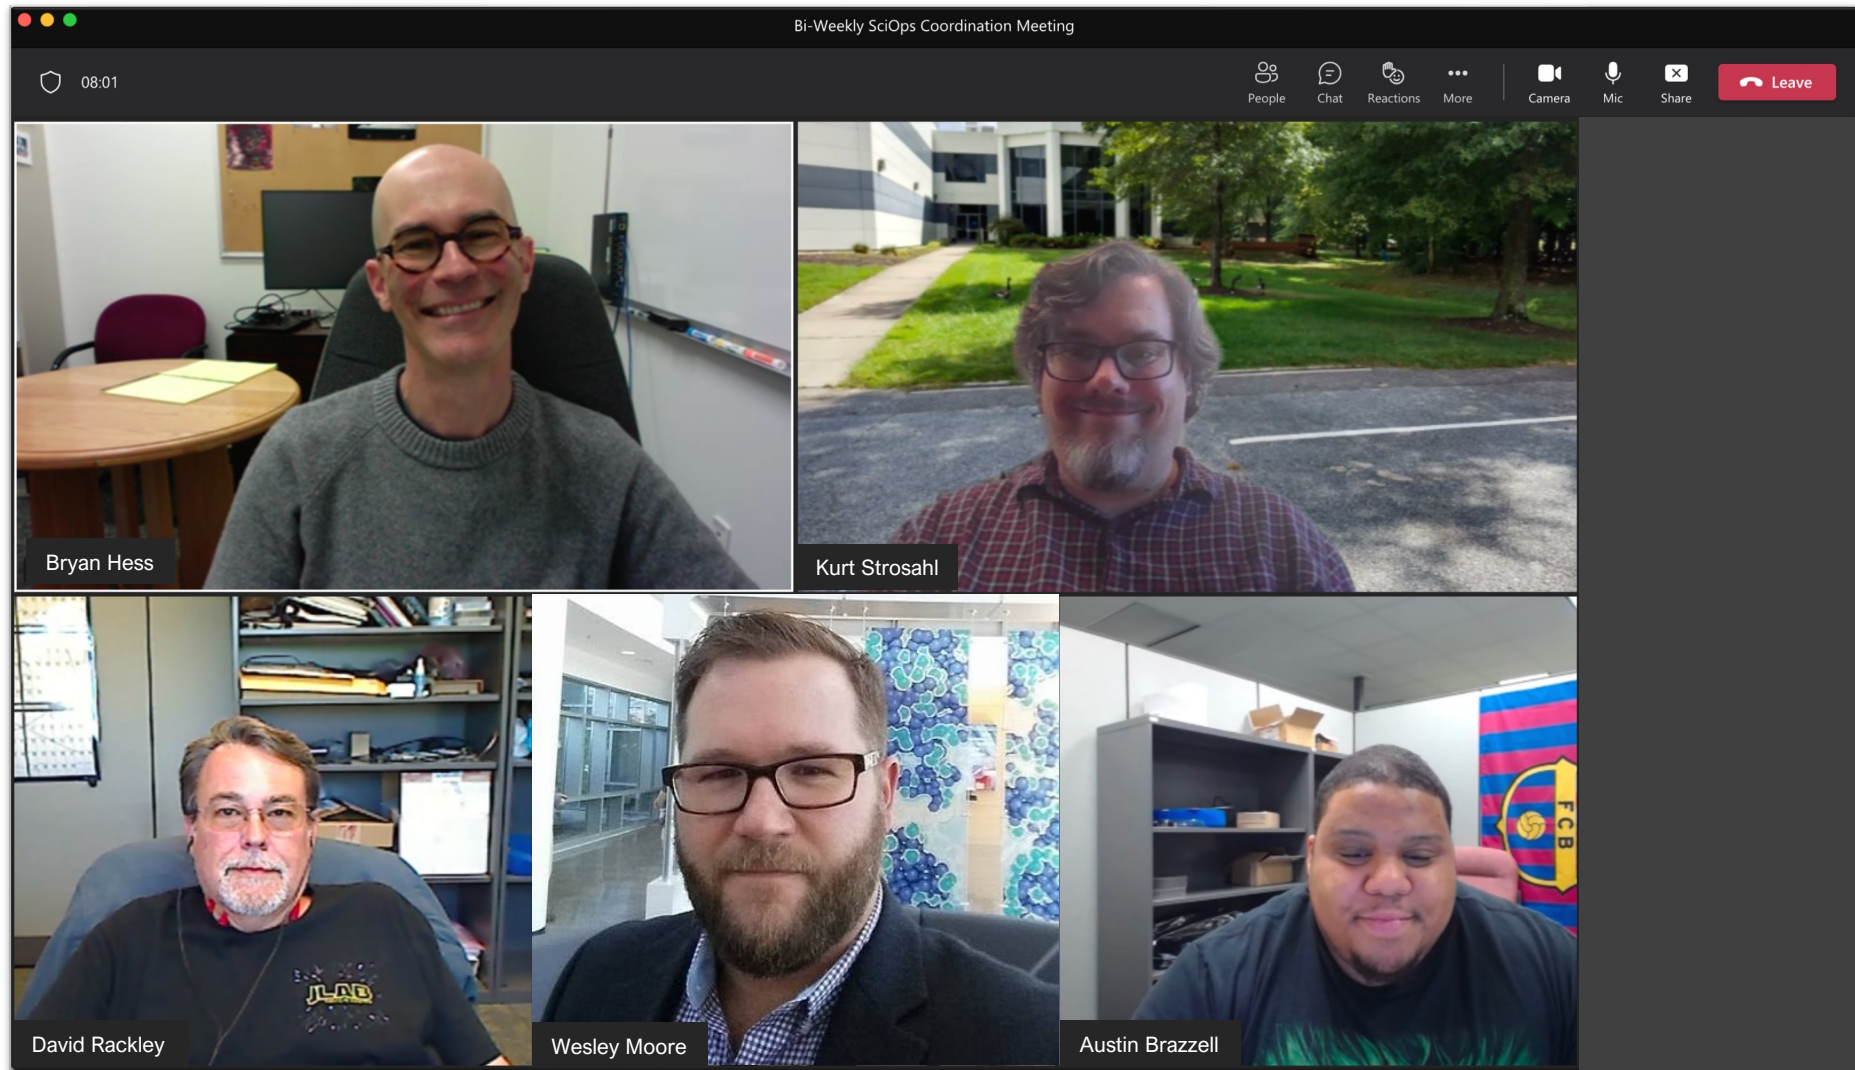


Jefferson Lab Organization – hiring and growing

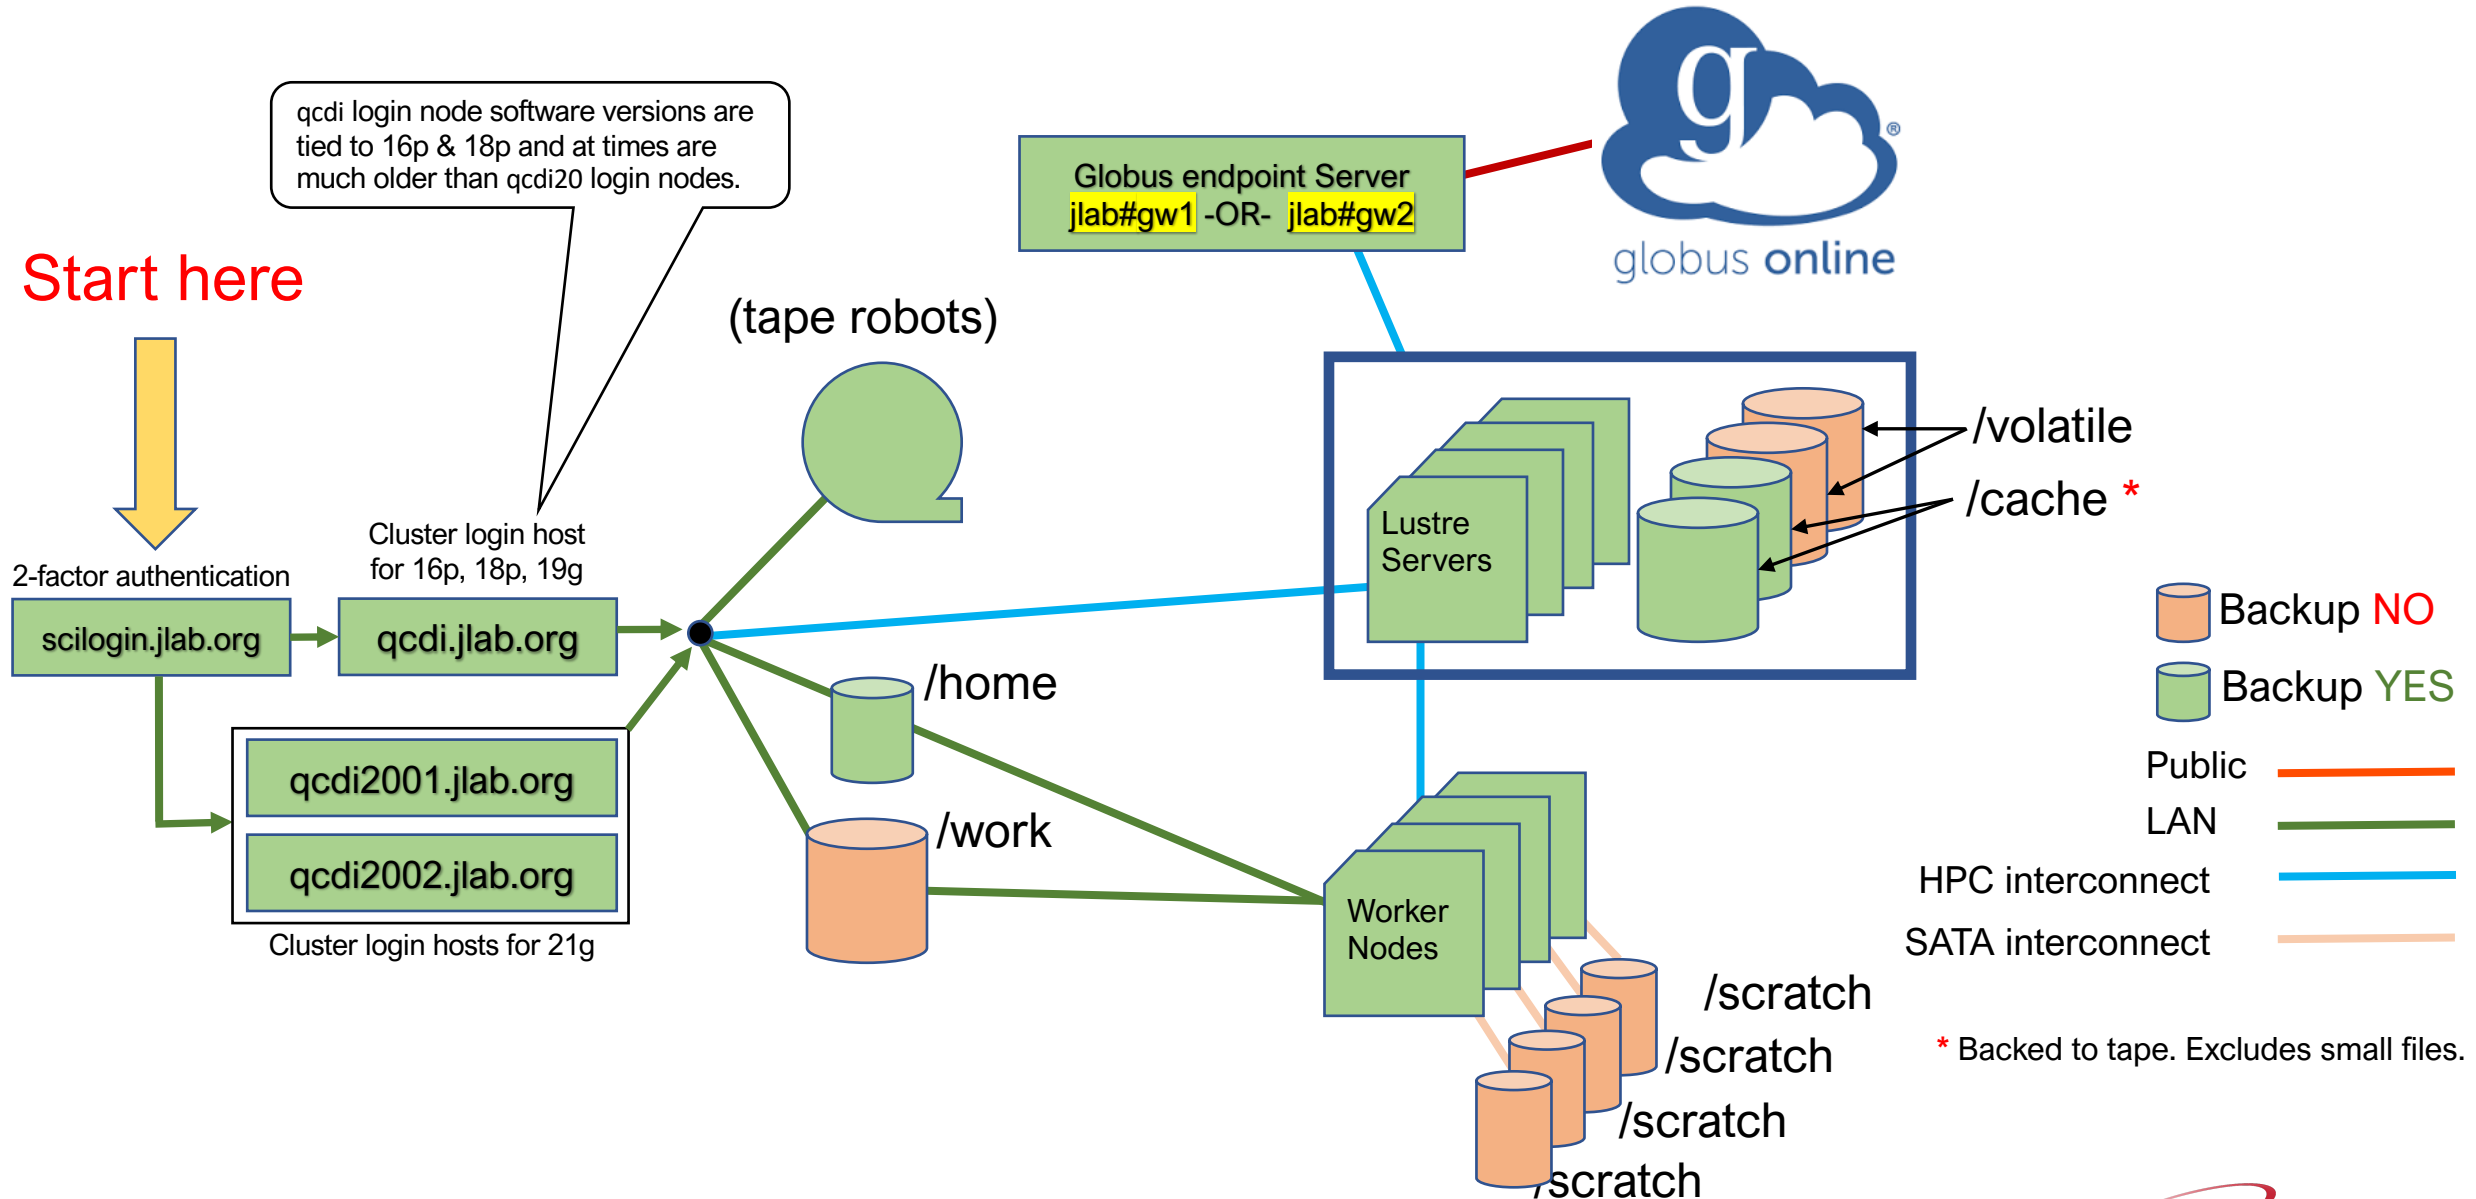


Existing Job Openings
[HTC hardware Specialist](#)
[Storage Hardware Lead](#)
[Infrastructure Developer](#)
[HPC-HTC Software Specialist](#)

Your JLab Lattice QCD Support Team



JLab Cluster Layout Summary



Current computing as pledged for 2023-24 USQCD allocations

- Compute

- 400 node Xeon Phi / KNL cluster (“16p/18p”)

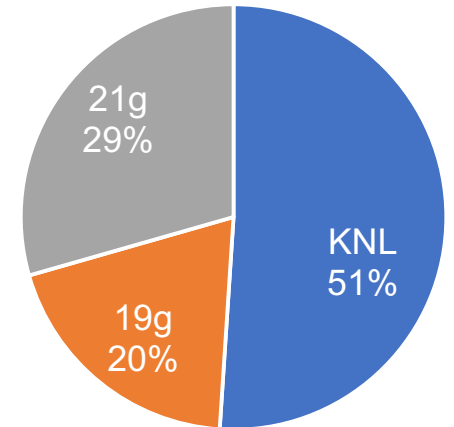
- Single socket 64 core KNL (with AVX-512 8 double / 16 single precision) 192 (98) GB main memory / node 16p (18p)
- 32GB high bandwidth on package memory (6x higher bandwidth)
- 100 Gbps bi-directional OmniPath network fabric (total 25GB/s/node) 32 nodes / switch, 16 up-links to core / switch
- Total: 93.3M Sky-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
- Total: 35.7M Sky-core-hours

- 8-node AMD GPU Cluster (“21g”)


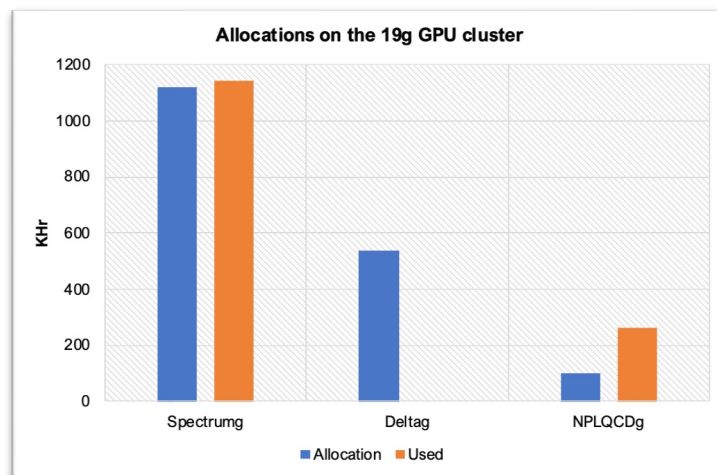
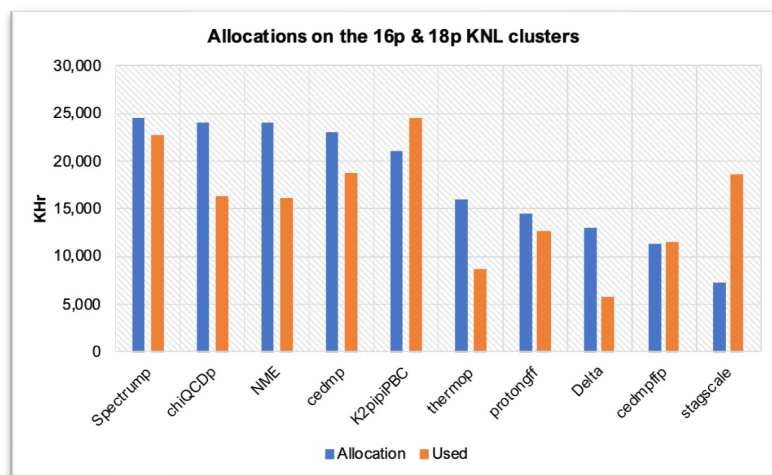
- Eight-CPU AMD MI100 nodes with Inter-GPU Infinity interconnect. 32GB memory per GPU, 1TB memory per node.
- Each on 100g InfiniBand Fabric
- Total: 53.8M Sky-core-hours



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.
 - These areas are now on two separate file servers. Thank you for your patience during the data migration
 - /work NFS on ZFS and is not backed up (<https://lqcd.jlab.org/lqcd/workDisk>)
 - /home is flash storage and is backed up
- Tape Storage
 - All writes are now to LTO8 (12TB) tape media
 - To date LQCD accumulated storage is 13.8PB on tape (<https://lqcd.jlab.org/lqcd/cacheDisk/project>)
 - 12PB on lattice-p "permanent"
 - 1.8PB 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

2022-2023 Allocations Summary



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

87%

80%

76%

We have setup a 21g_test queue for quick turnaround on 21g for code development and optimization jobs esp. to encourage new users.

User Documentation & how to ask for support

The screenshot shows the Jefferson Lab LQCD website interface. At the top right, there are navigation links: "Getting Started", "Staff Members", and "Support". Two red arrows point to these links. On the left sidebar, there is a "Documentation" section with a sub-menu containing "User's Guide", "Unrecoverable File", and "News Archive". A red arrow points to the "User's Guide" item. The main content area is titled "Jefferson Lab LQCD" and contains a mission statement, a reminder about SSH access (dated Mar-14-23), and a tip about the 21g cluster (dated Dec-14-21). Below this, there are two cluster status sections: "knl Cluster" and "gpu Cluster", each with three donut charts showing usage percentages. The "knl Cluster" charts show All (438), 18p (176), and 16p (262). The "gpu Cluster" charts show All (40), 21g (8), and 19g (32). Below the cluster status, there is a "File System" section with five cards showing usage: lustre (77%), /cache (104%), /volatile (87%), /work (47%), and /home (39%). At the bottom, there are two sections for Slurm jobs: "Slurm Outstanding jobs" (Running: 15, Pending: 451, Held: 0, Other: 0) and "Slurm past 24 Hrs finished jobs" (success: 148, Failed: 0, Timeout: 0, Cancelled: 0). At the bottom left, it says "Last updated: Tue Apr 18 2023 13:36:06 GMT-0400 (Eastern Daylight Time)".

<https://lqcd.jlab.org>

User Documentation in the form of Service NOW Knowledge base articles

The screenshot shows the LQCD User Portal interface. At the top left is the text "LQCD User Portal" and at the top right is a "Log in" link. Below the header is a search bar labeled "Search Knowledge" with a magnifying glass icon. The main content area is divided into two columns. The left column is titled "Categories" and lists: "Computing Resources" (1 article), "FAQ" (1 article), "New user accounts" (1 article), "Software" (3 articles), and "Storage Resources" (2 articles). The right column is titled "Most Viewed Articles" and lists: "New Account Request" (55 Views), "LQCD Hardware at Jefferson Lab" (54 Views), "Running Jobs on the JLab LQCD Clusters." (24 Views), "Frequently Asked Questions" (19 Views), and "Various filesystems available on the JLab LQCD clusters." (7 Views).

To help us improve the quality of our documentation You can rate or comment on each article !!



The feedback section includes a "Helpful?" poll with "Yes" and "No" buttons, showing "100% found this useful". To the right is a "Rate this article" section with five stars. Below these is a comment box with a placeholder "Click here to comment on this article..." and a grey circular icon containing the letter "G".

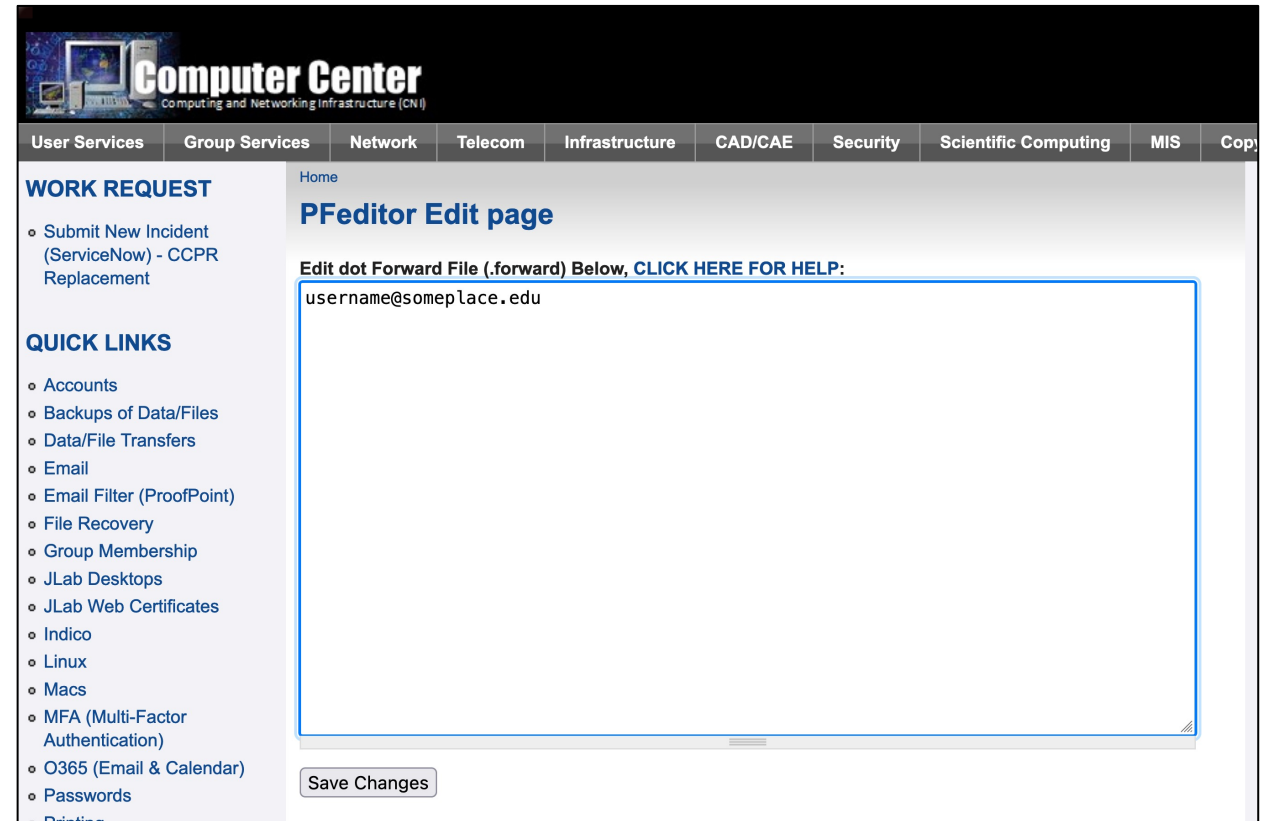
https://jlab.servicenowservices.com/lqcd?id=kb_view2

A few words from the operations team – 2022-23

- Globus transfer nodes were replaced with new hardware and software was upgraded to the latest.
- The Intel OneAPI compiler suite and HPC Toolkit now available. Added RHEL 8/Stream 8 support
- Since March 21st, MFA logins are now required (scilogin.jlab.org is the MFA gateway)
- Many 16p and 18p nodes have been rebuilt due to hard drive failures (lifetime issue)
- For 2022 we had 96 support tickets and 92% of the tickets were resolved within 3 days or less.
- Reminder
 - /volatile and /cache are on Lustre. /volatile is not backed up. /cache is written to tape.
 - Infiniband core network is EDR Infiniband. We will upgrade to an HDR core this FY and continue to support EDR and Omni Path.

Remember to forward your JLab email

- Email from the lab will be sent to your registered JLab email address so please make sure to check your JLab issued email.
- You can forward your JLab email to your preferred email account using this web page <https://cc.jlab.org/pfeditor/edit> (authentication required, use your JLab CUE account).
- Email forwarding is highly recommended, so you do not miss out on important announcements and communications related to your JLab accounts.



Sample email forwarding PFeditor page once you authenticate with your JLab CUE account

Questions?

| | |
|-------------------------------|---|
| JLab LQCD Main Page | https://lqcd.jlab.org/ |
| New User Accounts | https://jlab.servicenowservices.com/lqcd?id=kb_article_view&sysparm_article=KB0014813 |
| Submit a helpdesk ticket form | https://lqcd.jlab.org/lqcd/support |
| Knowledgebase articles | https://jlab.servicenowservices.com/lqcd?id=kb_view2 |
| FAQ's | https://jlab.servicenowservices.com/lqcd?id=kb_article&sysparm_article=KB0014827 |

