Fermilab Facilities Report

Smita Darmora

USQCD All-Hands Collaboration Meeting

April 20, 2023

The Computational Science and AI Directorate (CSAID)

- The Computational Science and AI Directorate (CSAID) provides software solutions and deploys and operates scientific computing facilities in support of the Fermilab program.
- Its staff works closely with the other laboratory organizations to deliver world-class computing services, operations and software engineering support to Fermilab-based experiments, CMS and the high-energy physics community at large
- It interacts with the experimental community to determine capacity, availability and capability requirements and to procure resources as necessary.
- It engages in R&D activities required to maintain or advance capabilities necessary for the success of Fermilab's future physics program.

James Amundson
Associate Lab Director, Computational Science and AI Directorate

https://computing.fnal.gov/organization/
https://computing.fnal.gov/computing-facilities-and-middleware/

The HPC Support Team

- Applications / Programming - - **Jim Simone** (Site Liaison), **Alexei Strelchenko** (Developer)
- Systems / Hardware - - Ed Simmonds (GL), **Tim Skirvin** (HPC Architect)
- Services / Support - - Ken Herner (GL & Site Manager), **Smita Darmora** (Site Manager)
- Batch Management - - **Farrukh Kahn** (GL)
- Storage Services - - Robert Illingworth (GL), **Dan Szkola** (Lustre/ZFS)

Fermilab LQCD support - Service Desk

- Emailing lqcd-admin@fnal.gov is the best way to reach us and will elicit the quickest response.
- LQCD-Admin staff use Service Now (SNOW) ticketing system for tracking incidents and requests.
- Email sent to lqcd-admin@fnal.gov automatically generates an Incident ticket.
  - Soon we will stop doing automatic tickets, because of so much spam.
- Ticket Types in SNOW (ServiceNow)
  - INC – Incidents are for when something is broken. Higher priority and tracked closely
  - RITM – Requests are for all other questions or assistance. Still tracked but less detailed reporting up to management or DOE
- We convert INC to RITM when appropriate.
- Your email reply becomes part of the worklog of the ticket number in the subject.
- Response time as well as time to resolve are tracked and reported.
- Please do not contact any of the Fermilab LQCD support staff directly.

Our goal is to resolve issue within 3 business days.

Subscribe to the lqcd-users@fnal.gov

To receive notifications about LQ operations and future announcements concerning the Fermilab lattice QCD clusters, please subscribe to the lqcd-users@fnal.gov mailing list.

To subscribe:

1. Send an e-mail message to listserv@fnal.gov.
2. Leave the subject line blank.
3. Type the following in the body of the email message: SUBSCRIBE lqcd-users FIRSTNAME LASTNAME

New User Accounts

To access USQCD computing resources at Fermilab users require the following three items:

- Visitor ID
- Kerberos account
- Unix accounts on the LQCD cluster machines

Visitor ID and Kerberos account

Submit an online application form for a Visitor ID and Kerberos account using this link.

Unix accounts on the LQCD cluster machines

After you have received an email with information about your Kerberos account, We need an email from the PI of the project verifying your affiliation to the project.

We are working on creating a ServiceNow (SNOW) form which will allow project PI's to automatically add people.

Fermilab's LQ1 Institutional cluster

- 181 nodes, 7,240 cores.
- Intel 6248 "Cascade Lake" 2.5 GHz) processors and an EDR Omni-Path fabric.
- 375 GFlops/node - - 59 TFlops total **
- Our /lustre1 filesystem is now 820 TB total. Current quotas add up to 742 TB
- Also watching quotas on /project (24 TB) and /home (1 TB) areas.
- Current Type A Allocations total 55.4 M sch (sky- core-hours)
- All Allocation Types total 57.94 M sch
- Delivered 2.9M Opportunistic sch so far, this PY

** not applying Amdahl's law.

All usage of the LQ1 cluster since July 1, 2022

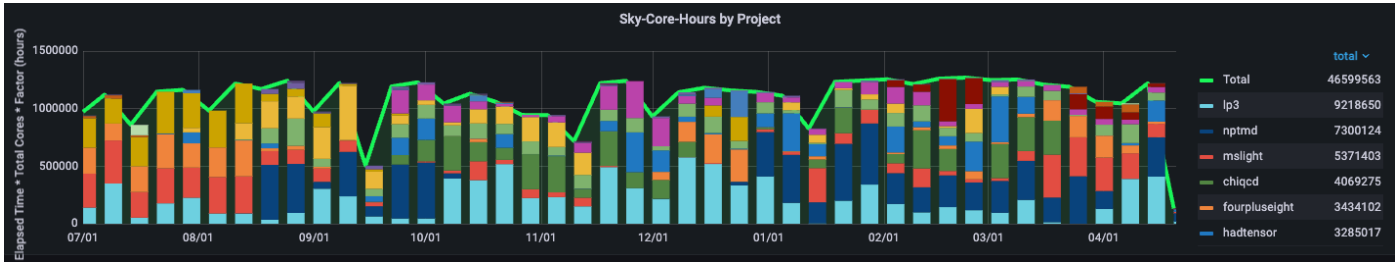

46599563.34

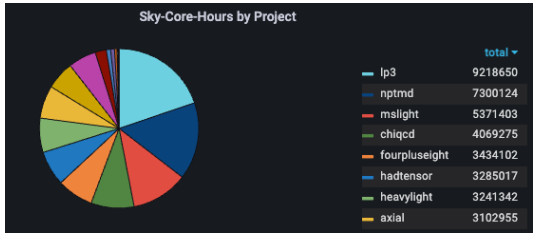

- Few dips in fully loading the cluster came when there were not enough jobs in queue to keep the cluster fully loaded.
- We have ability to make temporarily adjustment to keep the cluster busy.

Fermilab

Progress against allocations as of April 20, 2023

Project Name	Cluster	SPC Original Allocation (Sky-Core-Hours) ▲	Adjustments (Sky-Core-Hours)	SPC Adjusted Allocation (Sky-Core-Hours)	Project Used as of Jul 1, 2022 (Sky-Core-Hours)	Progress against Adjusted Allocation
lp3	FNAL-LQ1	13,700,000	-	13,700,000	9,224,209	67%
chiqcd	FNAL-LQ1	11,300,000	-4,171,925	7,128,075	4,157,553	58%
mslight	FNAL-LQ1	7,550,000	291,389	7,841,389	5,355,022	68%
axial	FNAL-LQ1	6,000,000	-355,078	5,644,922	3,105,684	55%
heavylight	FNAL-LQ1	6,000,000	-695,741	5,304,259	3,247,553	61%
nptmd	FNAL-LQ1	5,500,000	2,895,368	8,395,368	7,322,330	87%
ahisq	FNAL-LQ1	2,800,000	711,598	3,511,598	2,648,168	75%
hadtensor	FNAL-LQ1	2,400,000	1,324,389	3,724,389	3,283,634	88%

Not shown: the hours used by 5 Type B (1.9M), 3 Type C (26K), and 7 opportunistic (3.2M) projects

TOTAL	FNAL-LQ1	57,145,000	-	57,145,000	46,598,073	69.2%

These numbers from April 20 are at 80% through the program year

https://www.usqcd.org/fnal/clusterstatus/lq1/accounting.html

https://www.usqcd.org/jeopardy.pdf

Storage (disk and tape)

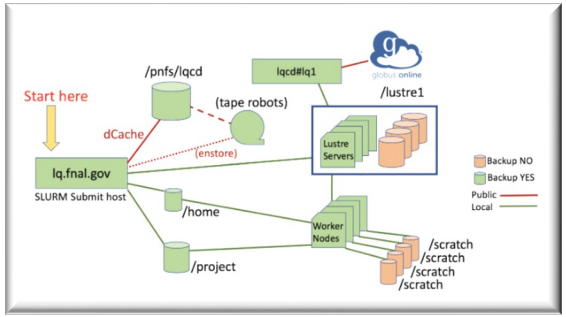

Area	Description
/project	Area typically used for approved projects. Visible on all cluster worker nodes via NFS file system. Backups nightly. Suitable for output logs, meson correlators, and other small data files. NOT suitable for fields e.g configs, quark propagators. Go to the Cluster Status page to check your /project disk usage.
/home	Home area. Backups nightly. Visible on all cluster worker nodes via NFS. Not suitable for configs or props. Can be used as a "run" directory for light production or testing. A quota of about 6 to 10 GB per home directory. Run `zfsquota` on `lq.fnal.gov` to check your home area disk usage.
/pnfs/lqcd	Enstore Tape storage. Visible on cluster login head nodes only. Ideal for permanent storage of parameter files and results. Must use special copy command: ' `dccp` '
/lustre1	Lustre storage. NO backups. Visible on all cluster worker nodes. Ideal for temporary storage (~month) of very large data files. NOT suitable for large number of small files. Disk space usage monitored and disk quotas enforced. Go to the Cluster Status page to check your project's /lustre1 disk usage.

File Transfer :

- Please use lqio.fnal.gov for large IO transfers to Lustre or tape.
  - Has 100GigE interface to offsite networks.
- Globus end point lqcd#lq1 for /lustre1.

- /scratch with 890 GB. No backups, suitable for writing all sorts of data created in a batch job.

https://computing.fnal.gov/lqcd/filesystems/

Fermilab

Quotas for allocated storage in Lustre and /project

- We track /lustre1 usage and quota based on group ownership.
- Periodically, we run scripts to correct group ownerships and sticky bits
- Since Lustre is designed for large file access, we have implemented limits on storage space and number of files

Low (>)	High (<=)	Number of files
0	5 TB	1.0 M
5 TB	50 TB	2.5 M
50 TB	100 TB	5.0 M
100 TB	150TB	7.5 M
150 TB		10.0 M

- We track /project area usage and quota based on ZFS project ids.
- Since the number of files in ZFS is not an issue, there are no limits

https://www.usqcd.org/fnal/clusterstatus/lq1/diskusage.html

Data Management Planning

- Each project is responsible for their own data management plan.
- You need to keep copies of critical data somewhere safe. **Lustre is not backed up.**

- A project's storage allocation is tied to the program year. Think of it like a lease on an apartment.
  - At the end of the year, you either renew your lease, negotiate to a larger or smaller space, or you vacate and move to another location.
  - Clear out promptly please. We need to provide the space to the next "tenant".

- Short term tape storage will be separate tapes by program year vs Long term.

https://computing.fnal.gov/lqcd/fermilab-data-management-guidelines-policies/

SLURM Limits and Dispatch Priority

- Maximum number of nodes per job and nodes per account is 64.
- Maximum number of jobs submitted per account is125.

Name	Description	Priority	GrpTRES	MaxWall	MaxJobsPU	MaxSubmitPA
admin	admin testing	600				
test	quick tests of scripts	500	cpu=80	00:30:00	1	3
normal	Normal QoS (default)	250				125
opp	unallocated/opportunistic	10		08:00:00		125

Job's requesting shorter wall-time limits are more likely to start sooner as Backfill jobs.

https://computing.fnal.gov/lqcd/slurm/ and https://computing.fnal.gov/lqcd/job-dispatch-explained/

Scheduled downtime

- Two downtimes per year for Kernel updates and other Security patches.
  - Early July during transition to new Allocations
  - Early January as the 6-month mark.
- We are considering more frequent reboots based on issues seen in Dec-Jan, where the nodes were running out of the memory.
- When we need to do other scheduled maintenance, it will typically be on the third Wednesday of the month.
- We will continue to use the lqcd-users@fnal.gov mailing list to announce both scheduled downtimes and unscheduled outages.
- The scheduled maintenance window for Enstore and dCache is the third Wednesday of each month, beginning at 8am, normally 4-8 hours.

Issues since last Fall (July 1, 2022)

Lustre issue fall 2022:

- One of the nodes failed. Lustre was inaccessible for a week.
- As far as we are aware, no data was lost.
- Affected machine was taken out and than it was subjected to intensive testing for few months and it never showed any problem after that.
- As a consequences when we got it back some users were showing over quota. even if they weren't.
- It has been more than two months since we have received any complaints regarding the over quota issue.
- **Issue resolved,** but we still don't know what the root cause was (can't reproduce). Very subtle problem.

Fermilab GPU cluster : LQ2

13 Jan 22 – Hardware requirements committee kickoff meeting.
28 Mar 23 – Purchase officially awarded to vendor (Koi).
01 Oct 23 – LQ2 cluster becomes an allocated USQCD resource.

Features:

- 18 nodes each with 4 NVIDIA A100 GPU
- Dual AMD EPYC 7003 Series Processor.
- 960 GD SSD scratch disk.

James Simone: https://indico.mit.edu/event/729/contributions/2033/

We asked you to use the lqcd-admin@fnal.gov email address, but if you have questions about this presentation, you are welcome to contact me directly:

Smita Darmora <sdarmora@fnal.gov>

Thank you

Three rainbows cross the sky over the Tevatron ring