



Fermilab Facilities Report

Ken Schumacher USQCD All-Hands Collaboration Meeting 1-2 May 2020

Introduction

- 1. Fermilab USQCD "Institutional" cluster.
- 2. Storage (disk and tape).
- 3. Progress against compute & disk allocations.
- 4. Allocation year 2019-20 (PY19-20) Job statistics.
- 5. SLURM tips and tricks.
- 6. Upcoming upgrades and changes.
- 7. Kerberos & SSH troubleshooting tips.
- 8. User support and web portal.
- 9. Questions?



Fermilab Scientific Computing Division

- The Scientific Computing Division (SCD) provides software solutions and deploys and operates scientific computing facilities in support of the Fermilab program.
- SCD interacts with the experimental community to determine capacity, availability and capability requirements and to procure resources, as necessary.
- SCD engages in R&D activities required to maintain or advance capabilities necessary for the success of Fermilab's future physics program.
- Active Archive Facility (190PB on tape), CMS Tier-1 Center (24k cores), Data storage and handling, Data Centers, FabrIc for Frontier Experiments (FIFE), General Purpose Grid, HEP Cloud and HPC.



Elizabeth Sexton-Kennedy Chief Information Officer



James Amundson Head of SCD



Fermilab USQCD "Institutional" cluster

- 183 nodes, 7,320 cores.
- Intel 6248 "Cascade Lake", 2:1 oversubscribed EDR Omni Path.
- FY19: Initial purchase, 112 nodes.
- FY20: Expansion purchase, 71 nodes.
- Expansion installation completed successfully following COVID-19 shelter-in-place restrictions.
- Currently 3 Fermilab projects run in "opportunistic" mode.
- Total cost: \$1.63M
- 325 GFlops/node
- 59 TFlops total **

** not applying Amdahl's law.





Storage (disk and tape)

- Please use lqio.fnal.gov for large IO transfers to Lustre or tape.
 - Has 100GigE interface to internet.
- Best effort to resolve Globus related issues.



Area	Description	Area	Description
/home	Home area. Visible on all cluster worker nodes via NFS. NOT suitable for configurations or propagators. SUITABLE as a "run" directory for light production or testing. Quota of 30 GB per user. Backups nightly.	/lustre1	Lustre storage. Visible on all cluster worker nodes. NOT suitable for large number of small files. SUITABLE for temporary storage (~month) of very large data files. Disk space usage monitored, and disk quotas enforced. NO backups.
/project	Area typically used for approved projects. Visible on all cluster worker nodes via NFS file-system. NOT suitable for fields e.g. configurations, quark propagators. SUITABLE for output logs, meson correlators and other small data files. Backups nightly.	/pnfs/lqcd	Tape storage. Visible on login head nodes only. NOT suitable for large number of small files, compression highly recommended. SUITABLE for permanent storage of parameter files and results. Must use special copy command: 'dccp'

https://www.usqcd.org/fnal/v2/filesystems.html



Progress against compute allocations as of April 28, 2020

Project Name	Cluster	SPC Original Allocation (Sky-Core-Hours)	Adjustments (Sky-Core-Hours)	SPC Adjusted Allocation (Sky-Core-Hours)	Project Used since Jul 1, 2019 (Sky-Core-Hours) ▲	Progress against Adjusted Allocation	Remaining Allocation (Sky-Core-Hours)	30-day usage as of 04/28/2020	30-day burn rate as of 04/28/2020 (% of Alloc.)	Annual Pace	
chiqcd	FNAL-LQ1	7,000,000	1,400,000	8,400,000	8,499,143	101%	0	1,258,168	15.0%	Apr-24-2020	Π
vcbok	FNAL-LQ1	7,000,000	1,400,000	8,400,000	7,407,428	88%	992,572	2,117,254	25.2%	Jun-7-2020	11
fourpluseight	FNAL-LQ1	2,500,000	857,822	3,357,822	3,297,737	98%	60,085	226,614	6.7%	May-3-2020	1
rhgbbar	FNAL-LQ1	3,500,000	422,781	3,922,781	2,975,781	76%	947,000	414,353	10.6%	Aug-2-2020	1
safe	FNAL-LQ1	100,000	-	100,000	2,451,389	2451%	0	442,059	442.1%	Jul-13-2019	Г
lqcdadmin	FNAL-LQ1	-	-	-	1,816,659	-	-	8	-	-	1
axial	FNAL-LQ1	1,000,000	119,397	1,119,397	1,355,404	121%	0	564,568	50.4%	Mar-6-2020	1
heavylight	FNAL-LQ1	-	-	-	1,157,262	-	-	-	-	-	1
4fermi	FNAL-LQ1	100,000	-	100,000	323,726	324%	0	322,663	322.7%	Oct-2-2019	1
qis_algo	FNAL-LQ1	-	-	-	90,368	-	-	-	-	-	1
stgmugm2	FNAL-LQ1	100,000	3,400,000	3,500,000	73,159	2%	3,426,841	29,647	0.8%	Jan-19-2059	Π
qcdloop	FNAL-LQ1	-	-	-	30,078	-	-	21,358	-	-	Г
g4	FNAL-LQ1	-	-	-	23,516	-	-	13,212	-	-]
hadtensor	FNAL-LQ1	100,000	-	100,000	1,421	1%	98,579	1,224	1.2%	Sept-6-2077]
lp3	FNAL-LQ1	100,000	-	100,000	425	0%	99,575	425	0.4%	Jan-18-2214]
osgcovid19	FNAL-LQ1	-	-	-	18	-	-	18	-	-]
hiq2ff	FNAL-LQ1	100,000	-	100,000	-	-	100,000	-	-	-	1
mslight	FNAL-LQ1	100,000	-	100,000	-	-	100,000	-	-	-	1
higgs	FNAL-LQ1	100,000	-	100,000	-	-	100,000	-	-	-]
s1080	FNAL-LQ1	100,000	-	100,000	-	-	100,000	-	-	-]
nptmd	FNAL-LQ1	20,000	-	20,000	-	-	20,000	-	-	-]
TOTAL	FNAL-LQ1	21,920,000	7,600,000	29,520,000	29,503,514	-	6,044,652	5,411,571	-	-	Π

https://www.usqcd.org/fnal/clusterstatus/lg1/accounting.html



Progress against disk allocations

Project Name	Filesystem	used (TB)	quota (TB)	limit (TB)	grace	files ▲	quota	limit	grace
fourpluseight	/lustre1	17.04	27.0	27.0	-	23,136,335	0	0	-
rhqbbar	/lustre1	79.09	100.0	100.0	-	8,435,533	0	0	-
lqcd	/lustre1	63.39	65.0	65.0	-	8,330,341	0	0	-
g4	/lustre1	5.87	10.0	10.0	-	808,877	0	0	-
higgs	/lustre1	23.00	23.0	23.0	-	673,430	0	0	-
vcbok	/lustre1	3.68	15.0	15.0	-	353,997	0	0	-
4fermi	/lustre1	0.03	1.0	1.0	-	344,615	0	0	-
lp3	/lustre1	0.31	2.0	2.0	-	175,412	0	0	-
heavylight	/lustre1	27.96	35.0	35.0	-	164,626	0	0	-
nptmd	/lustre1	0.00	1.0	1.0	-	93,015	0	0	-
axial	/lustre1	0.43	25.0	25.0	-	53,331	0	0	-
qis_algo	/lustre1	0.02	0.2	0.2	-	6,116	0	0	-
chiqcd	/lustre1	27.60	100.0	100.0	-	916	0	0	-
hadtensor	/lustre1	0.88	1.0	1.0	-	409	0	0	-
stgmugm2	/lustre1	3.24	20.0	20.0	-	228	0	0	-
qcdloop	/lustre1	0.00	0.2	0.2	-	2	0	0	-
hiq2ff	/lustre1	0.00	1.0	1.0	-	2	0	0	-
safe	/lustre1	0.00	1.0	1.0	-	1	0	0	-
\$1080	/lustre1	0.00	1.0	1.0	-	1	0	0	-
mslight	/lustre1	0.00	1.0	1.0	-	1	0	0	-
TOTAL	-	252.54	429.4	429.4	-	42,577,188	-	-	-



- Smallest file: 0 bytes
- Oldest file: 02-14-2012 (access time)
- Largest file: 950GB
- Each project should have a plan for managing their data over its lifetime.

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

7

Allocation year 2019-20 (PY19-20) Job statistics





SLURM Tips and Tricks

- Maximum number of nodes per job that can be requested is 88.
- Maximum number of jobs per project that can be submitted is 125.
- Job's requesting shorter wall-time limits are more likely to start sooner as Backfill jobs.

QoS Name	QoS Description	Priority	Group Resource Limit	Maximum Wall-time	Maximum Jobs Per User	Maximum Job Submits Per Account
admin	Reserved for Administrators	600				
test	Quick turnaround testing	500	cpu=80 (node=2)	00:30:00	1	3
normal	Normal QoS (default)	300				125
орр	Opportunistic	0		08:00:00		125

https://www.usqcd.org/fnal/v2/slurm.html and https://www.usqcd.org/fnal/v2/slurmdispatchpr.html

SLURM Tips and Tricks (cont'd)

- Launching MPI processes
 - Near the beginning of your batch script, prior to launching an MPI process you should ensure software modules required by the batch script have been loaded, for example if using gnu8 and openmpi:

module purge

module load gnu8 openmpi3

• We recommend using **srun** to launch all three MPI implementations since it is then possible to use command line options to specify the distribution and binding of MPI ranks to CPU cores and local memory. Reasonable affinity choices are:

--distribution=cyclic:cyclic --cpu_bind=sockets --mem_bind=no

https://www.usqcd.org/fnal/v2/software.html

Upcoming upgrades and changes

- LQCD resources opened for OSG COVID-19 research.
 - On LQ1 'osgcovid19' is setup as an "opportunistic" account.
 - On WC (was pi0 and pi0g) are exclusively 'osgcovid19'.
- Adding ~200TB to /lustre1. 610TB --> 810TB
- Dedicated NFS-server for /project and /home areas. Possible 1-day outage. We may be able to increase existing /project quotas.
- Currently USQCD users have access to all 183 LQ1 nodes. Project is paying for cycles on 171 nodes. Possible changes:
 - Additional Fermilab accounts may be added for "opportunistic" access to LQ1 nodes.
 - HEPCloud "Science Gateway" integration.



Kerberos & SSH troubleshooting tips

- Fermilab computer security policy dictates that all machines that allow remote logins from the internet do so using strong authentication and encryption. We use Kerberos for authentication and SSH for encrypted remote login sessions to the cluster server machines such as lq.fnal.gov and lqio.fnal.gov.
- A few common issues we have noticed users experience when using Kerberos:
 - kinit: krb5_get_init_creds: Error from KDC: CLIENT EXPIRED
 - What is my 5-digit Fermilab ID?
 - kinit: Preauthentication failed while getting initial credentials
 - kinit: krb5_get_init_creds: unable to reach any KDC in realm FNAL.GOV, tried 0 KDCs
- A few common issues we have noticed users experience when using SSH:
 - Permission denied (gssapi-keyex,gssapi-with-mic)
 - ssh: connect to host lq.fnal.gov port 22: Connection timed out

https://www.usqcd.org/fnal/v2/troubleshooting.html

User support and web portal

- <u>http://www.usqcd.org/fnal/</u> There is so much information available here
- Emailing <u>lqcd-admin@fnal.gov</u> is the best way to reach us and will elicit the quickest response. Please <u>do not</u> contact any of the Fermilab LQCD support staff directly. Emailing lqcd-admin will automatically open a support ticket which is tracked, and we will do our best to resolve your issue within 3 business days.
- Ticket Types
 - 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 not as much reporting up to management or DOE
- Emails generate INC which we will convert to RITM when appropriate.

https://www.usqcd.org/fnal/v2/contact.html

Questions ?

