



Fermilab LQCD Facilities Roadmap

Jim Simone

March 2, 2023

LQCD Facility Roadmap Issues

- hardware funding plans
- resource plans
- LQ2 GPU cluster procurement update
- LQ2 deployment planning

LQCD funding

- The DOE/SC HEP office provides funds to obtain LQCD computing resources from the labs (BNL, FNAL) through the LQCD-ext hardware initiative.
- The NP office also funds computing for LQCD at JLab.
- Compute resources at the three labs are allocated through the USQCD collaboration to researchers in both HEP and NP to maximize the science output of USQCD.
- LQCD-ext, USQCD, and the labs coordinate on compute resource planning to best match compute facilities to the USQCD science needs.
- Both BNL (CPU+GPU) and FNAL (GPU) started FY22/23 procurements.
- USQCD, LQCD-ext, and DOE/SC agreed only FNAL procurement would continue.
- FY23 BNL funds came to FNAL increasing the hardware budget ~50%.
- BNL and FNAL will alternate procurements – fewer but bigger procurements!

FNAL compute resource timeline

- Existing CPU cluster “LQ1”
 - LQCD-ext has subscribed all resources in FY23
 - Plan is for USQCD to continue to use LQ1 at least through FY24
- New GPU cluster “LQ2”
 - Joins LQ1 in the LQ complex, shared batch, storage, login servers
 - Deploy in Spring 23, friendly user testing through end FY23.
 - Comes fully under USQCD allocation process in FY24.
 - An allocated USQCD resource in the USQCD program year beginning July 1, 2024.
- Future clusters
 - Next procurement at FNAL would happen in FY25.
 - Architecture TBD: CPU+HBM, converged CPU+HBM+GPU, ??

LQ2 procurement status

The procurement is not final – some details remain confidential

Purchase recommendation

- Eighteen GPU workers designed for HPC computing:
- Four NVIDIA A100 each with 80 GB of HBM2e memory
- GPUs interconnected in a DGX4 “Redstone” topology
- dual AMD 32-core 2.8 GHz Milan CPUs
- 1 TB of RAM
- InfiniBand interconnect providing 400 Gbps b/w per worker

LQ2 performance will be about 2.7x the performance of 175 worker LQ1 CPU cluster for LQCD.

Timeline

- Aug 5 – draft RFP sent to procurement
- Oct 5 – procurement releases RFP to vendors on SAM
- Dec 28 – extended deadline for all bid submission materials
- Jan 9 – procurement okays technical evaluation of bids
- Jan 30 – procurement reveals purchase costs of bids
- Feb 1 – draft purchase recommendation set to procurement
- Mar 1 – working with procurement to finalize the award

LQ2 Deployment Tasks – Much work needs to be done

Task title	Assigned to	Priority	Due ↓	Bucket	
+ Add a task					
<input type="radio"/> SLURM configuration refactoring	FK Farrukh A. Khan		2/20	Blocked by depe...	...
<input type="radio"/> Award purchase to vendor	0/2 GC TS DF JS		3/10	To do	...
<input type="radio"/> HPC software deployment	0/4 JS AS		3/10	Work in progress	...
<input type="radio"/> Deliver and install hardware	1/1 GC DF		3/17	To do	...
<input type="radio"/> Hardware validation testing	0/2		3/22	To do	...
<input type="radio"/> LNET router deployment	0/5 GC DF		4/14	To do	...
<input type="radio"/> Resource usage reporting	0/6 JB Joseph B Boyd		4/14	To do	...
<input type="radio"/> LQCD project and user management	0/4 SD JB		4/14	To do	...
<input type="radio"/> Server configuration management	0/4 TS DF		5/4	Work in progress	...
<input type="radio"/> Lustre work	0/3 DS Daniel Szkola			To do	...
<input type="radio"/> Revise LQ documentation	0/5			To do	...