Fermilab 2022 GPU Cluster Acquisition

James Simone (HPC coordinator) and
Tim Skirvin (HPC architect)

April 2022

FERMILAB-SLIDES-22-023-SCD

Process timeline

- LQCD-ext plan calls for new compute resources at FNAL to be on-line early 2023.
- USQCD EC and SPC chairs, the FNAL HPC coordinator (J.S.) and LQCD-ext (Jo) meet and prefer a GPU cluster.
- ★ FNAL management agrees to host a GPU institutional cluster and charges an Acquisition Committee of domain experts from Lattice, HEP+CMS, Neutrino experiments, and Quantum Science with developing technical recommendations.
- The recommendations will be sent to the FNAL/SCD Procurement Committee to write the vendor RFP and conducts the purchase.

(Some!) Procurement challenges

- Supply Chain – Ongoing world-wide delays have already delayed the process and will specifically hurt interconnect and GPU timelines.
- Schedule Delay – The procurement timeline is very tight. The buy must be awarded to a vendor before the end of FY22.
- Budget – We will design an institutional cluster with input from CMS + HEP, the neutrino experiments, and Quantum Science. IC model: pool budgets to build something better. Funding commitment, however, is currently limited to just lattice QCD sources.
- Scheduling / Utilization – A relatively small but powerful cluster. Prolonged downtime of even a single worker is costly. ”Edge effects” makes job scheduling challenging on a small system.

Acquisition Committee

- Members: Peter Boyle, Chulwoo Jung, Thomas Junk (Neutrino), Dave Mason (CMS), Kevin Pedro (CMS), Gabe Perdue (Quantum + Neutrino), J. Simone (chair), Amitoj Singh, Tim Skirvin (co-chair), Frank Winter, and Mike Wagman.
- Formulate requirements, ”feeds and speeds”, etc, needed to meet the computational needs of the stakeholders.
- Identify potential benchmarks to be used to grade system performance. Make recommendations to the procurement committee.
- Examine vendor offerings and develop technical recommendations from the alternatives.
- Examine budget and spending scenarios for the new cluster.
- Report work to USQCD governance and Fermilab CD management. Transmit recommendations to the CD procurement committee.

Role of benchmarks

- Specifying benchmarks is an important part of the procurement process as well as rating the performance of existing USQCD resources.
- Many lattice QCD applications are good candidates for benchmarking.
- An impediment is a lack of good documentation and scripts covering steps need to build and run the code.
- Containerized pre-built benchmarks can eliminate the built steps. See specifications for MILC and Grid using NVIDIA's HPC Container maker tool.
- A community tabulation of benchmark results on different hardware would be very helpful!

The Fermilab LQ cluster complex

- The new GPU cluster ”LQ2” will be integrated with the Fermilab LQ1 cluster.
- Convenient for users.
- Cost savings through sharing of existing LQCD services.
- LQ2 will share login servers and slurm batch system with LQ1.
- Shared `/home`, Lustre, and `/project` file systems.

Technical recommendations

GPU Model:	NVIDIA A100-80 SXM	AMD MI250
GPU HBM2e Memory	80 GB	128 GB
GPU Interconnect	NVlink mesh	Infinity Fabric mesh
GPU Count	4 per node	
GPU-CPU Connect	PCIe Gen4 or better x16 lanes	
CPU Architecture	Intel Xeon or AMD EPYC	
CPU Cores	≥ 32	
System RAM	$\geq 1 TB$	
User Scratch Space	$\geq$ 1 TB local SSD or disk	

Interconnect to provive 400 Gbps BW per worker from either HDR InfiniBand (200 Gbps per port) or OmniPath (100 Gbps per port).

Recommended GPU connection topology

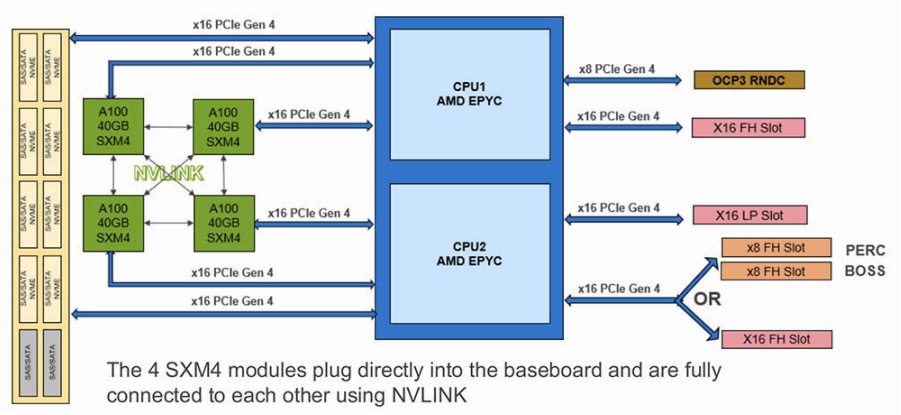

Example: Dell XE8545. Four A100 GPUs are interconnected by a point-to-point 600 GB/s NVlink mesh. Here each GPU is connected via PCIe Gen4 (x16) links to a CPU.

Considering the GPU options

AMD MI250	NVIDIA A100
+ Strategic choice to foster pricing pressure.	+ The default choice and market leader.
+ Good performance specifications on paper.	+ Good performance specs.
- LQCD frameworks are being ported, tested and tuned.	+ Frameworks originally developed on NVIDIA.
- Few HEP applications proven to run well.	+ HEP development mostly on NVIDIA.
- Learning curve for users.	+ Familiar to users.
- New operational experience.	+ Long experience.

Might AMD close the gap through very aggressive pricing combined with excellent performance on all benchmarks?

Procurement will not be just best price/performance but ”Best value with tradeoffs” scoring.