

**FY16 Acquisition Review Committee Report
for the
SC Lattice QCD Computing Project Extension II
(LQCD-ext II)**

Unique Project (Investment) Identifier: 019-20-01-21-02-1032-00

Operated at
Brookhaven National Laboratory
Fermi National Accelerator Laboratory
Thomas Jefferson National Accelerator Facility

for the
U.S. Department of Energy
Office of Science
Offices of High Energy and Nuclear Physics

Version 1.0

October 6, 2016

PREPARED BY:

The FY16 Acquisition Review Committee
*Rob Kennedy, Carleton Detar, Steve Gottlieb, Don Holmgren, Balint Joo, Shigeki Misawa,
Bob Mawhinney, James Osborn, Amitoj Singh*

FY15 LQCD-ext II User Survey Report Change Log

Version	Description	Date
0.5	Initial draft for review by Committee	7/12/2016
0.9	Revised draft ready for public consumption	8/1/2016
1.0	Final report	10/6/2016

Table of Contents

1	Introduction and Background.....	3
2	Deliverable: Recommendation to the Project Manager	3
3	Deliverable: Addressing The Charge Elements	3
4	Suggestions for Future Acquisition Reviews.....	6
5	Appendix: Charge to the FY16 Acquisition Review Committee	7

1 Introduction and Background

The purpose of the FY16 Acquisition Review Committee is to review and consider the proposed computing hardware acquisition plan for FY16 and provide input to the LQCD-ext II Project Manager regarding the alignment of the proposed procurement with the anticipated computing needs of the scientific program. The intent is to help ensure that the project is making the most effective use of project resources to further the USQCD scientific program.

Throughout this review, a variety of vendor NDA terms were in effect, different terms for different committee members, and business confidentiality for the TJNAF procurement was in effect. This severely hampered the desire by the committee to discuss quantitative metrics and actual pricing in order to provide the highest quality recommendation for the USQCD program. By agreement with vendors to address the immediate need for a preliminary recommendation, this committee report does not contain quantitative comparisons.

To be clear, this committee was charged to review the alternatives for their effectiveness in the USQCD program, not to reproduce the procurement process at TJNAF in recommending a best value alternative.

2 Deliverable: Recommendation to the Project Manager

The following recommendation was endorsed unanimously by the FY16 Acquisition Review Committee on May 31, 2016:

The preliminary recommendation of the Acquisition Review Committee is to proceed with the Acquisition Plan, given the information that we have.

3 Deliverable: Addressing The Charge Elements

Element 1: The near- and long-term demand (at a high level) for each hardware architecture in the existing portfolio and how the proposed acquisition will augment or complement the existing hardware portfolio;

- Near-term Demand:
 - According to the Report of the SPC (Anna Hasenfratz, 2016 All Hands Meeting), among type-A proposals, there was oversubscription rate (Request/Available) of:
 - BNL BG/Q 143%
 - Conventional 312%
 - Accelerated 120%
 - New FY16 Not included in the above categories.
- Long-term Demand: In the long-term we see:
 - Continued demand for CPU (whether conventional or MIC) technologies
 - Continued demand for GPU technology
 - A significant fraction of Conventional CPU demand metamorphosing into MIC CPU demand
 - Two technologies appear likely to be the most cost-effective for LQCD going forward:

- MIC (example today is KNL)
 - GPU (example today is Pascal)
- How will architecture augment or complement existing hardware portfolio:
 - We have test results that show un-optimized LQCD code can not only run on KNL, but run roughly cost-effectively compared to Broadwell, based on RFI cost estimates.
 - KNL can be effective for a larger fraction of USQCD applications than first thought, based on single KNL node tests.
 - Auto-vectorizable applications, those which use appropriate data structures, will not require explicit optimization to benefit from KNL. Some codes may benefit, others not.
 - KNL has tremendous cost-effectiveness potential for optimized code.
 - A proposed plan to eventually host a 256-node KNL cluster at JLab will complement the large memory footprint cluster (Pi0) at FNAL for conventional computing, where near-term demand is greatest by far, and provide an increasingly cost-effective platform as scientists optimize applications for the KNL features.
 - We conclude that, from what we know and see at this time, the potential for KNL is worth the risks associated with putting a new technology into production. It is reasonable to consider KNL as a procurement option.

Element 2: The availability of production software for use by enough of the USQCD collaboration to effectively utilize the capabilities of the proposed acquisition;

- At least 25% of the type-A conventional requests claimed that they could run on KNL according to the SPC survey subsequent to the Call for Proposals. Some committee members felt this was optimistic though given the context of the question. PI's were bound to be optimistic in case they could exploit the new technology in the future. Some large requests clearly could use KNL immediately, however, and those large requests would free up time on conventional resources that others with less compatible code could then use.
- We have test results that show un-optimized LQCD code can not only run on KNL, but run roughly cost-effectively compared to Broadwell, based on RFI cost estimates. This requires a boot-time node configuration that leaves the node less effective for highly optimized applications. If needed, a subset of a new KNL machine to be configured this way, reserved for un-optimized applications.

Element 3: The ability of the proposed acquisition, along with the existing hardware portfolio, to meet the established time-based performance goals for the computing project;

- Note: Performance goals for only FY17 – FY19 are affected by the FY16 Acquisition since the new machine is not in production use until very late FY16.
- With existing portfolio, we are optimistic about meeting performance goals with the KNL option.

- Focus was on KNL since near-term demand oversubscription is higher for conventional. The addition of BNL IC GPU's should help address GPU demand despite some GPU retirements.
- We do not see anything that will prevent us from meeting established performance goals that we cannot handle with existing plans.

Element 4: The capability of the project team to effectively support the computing hardware in the proposed acquisition, in terms of 1) meeting system uptime target goals; and 2) supporting the user community in the use of the newly acquired hardware;

- We acknowledge that new silicon has historically had issues. We cannot explore this however beyond committee members' experience on pilot/development machines:
 - i. From what one committee member could see, it's a standard Linux machine with standard tools.
 - ii. System can be painful when it arrives since it lacks some tools and utilities. However, a KNL cluster would use a standard rebuilt environment.
- There are risks, but from our limited experience, we do not see a serious issue in the areas of (1) or (2).

Element 5: The alignment of the computing hardware in the existing portfolio and new acquisition with vendor technology roadmaps; and with the technology roadmaps of leadership-class facilities at which USQCD collaboration members run scientific software codes.

- KNL is a major element of Intel's roadmap. We do expect to see some feature mixing and merge between conventional Xeon and Xeon Phi line over time.
 - KNL will be used by NERSC's CORI Phase II, Argonne LCF Theta.
 - KNH (Knight's Hill, successor to KNL) will be used by Argonne LCF Aurora.
- Broadwell was not discussed in detail except as a reference mark for KNL.
- NVIDIA Pascal GPU's are also in vendor and LCF roadmaps, but will arrive too late for the FY16 LQCD acquisition.
 - BNL IC allocation will maintain the total GPU performance level for now as older GPU clusters at FNAL and JLab are retired. However, the BNL IC will feature NVIDIA K80 accelerators, not Pascal GPU's.

4 Suggestions for Future Acquisition Reviews

These suggestions are not listed in any particular order:

Suggestion 1: Overall, we felt the Acquisition Review Committee was a useful tool in the Acquisition Process. We encourage that it become a standard part of the LQCD acquisition process. It helped foster communication between LQCD Site Architects and Project Office, and USQCD, in particular SPC, members that led to a better understanding of the technology choices.

Suggestion 2: We encourage a clearer statement of the review charge to clarify its role within the LQCD acquisition process. There was some confusion about the LQCD acquisition process due to words used in the charge that differed from acquisition process documentation.

Suggestion 3: We encourage establishing future Acquisition Review Committees much earlier in the acquisition process.

- This will allow the committee to provide input to the actual procurement process, such as which benchmarks to use in the best-value assessment. Some committee members would like to have seen more USQCD-specific benchmarks involved in the procurement.
- This may allow some non-acquisition-host-site members to participate directly in the procurement process, especially Site Architects, given enough time to gain clearance and be accepted by the procurement department at the host site.
- This may allow time to address NDA constraints more directly to allow freer communication within the committee of quantitative performance metrics.

Suggestion 4: The prioritization model presented by Chip Watson at the AHM enabled a fruitful high-level discussion without going into deep technical details. If this could be done earlier, then we could engage more of the community at a high level in time to have a large and non-disruptive impact on the Acquisition Plan and procurement.

Suggestion 5: We encourage tests of single-rail versus dual-rail on a future KNL acquisition to determine the impact on inter-node bandwidth. This should be one of the options considered for the FY17 acquisition: to purchase fewer nodes and outfit all nodes for dual-rail.

Suggestion 6: We propose that this committee reconvene at some point after the FY16 Acquisition machine is in service to re-evaluate some of the charge elements with some experience of running USQCD applications on the actual production cluster.

5 Appendix: Charge to the FY16 Acquisition Review Committee

LQCD-ext II FY16 Acquisition Review Committee Charge

April 13, 2016

Purpose

On an annual basis, the LQCD-ext II Computing Project typically executes one or more large purchases of computing hardware to augment the existing hardware portfolio operated by the project. The computing hardware operated and managed by the project is used by the USQCD Collaboration in support of its scientific program. In 2016, project funds will be used to purchase and deploy new computing hardware at Jefferson Lab (JLab). Following standard project practices, the FY16 Acquisition Plan has been prepared and is being implemented by Chip Watson, the LQCD-ext II Site Architect and Site Manager for JLab.

The purpose of this committee is to review and consider the proposed computing hardware acquisition plan for FY16 and provide input to the LQCD-ext II Project Manager regarding the alignment of the proposed procurement with the anticipated computing needs of the scientific program. The intent is to help ensure that the project is making the most effective use of project resources to further the USQCD scientific program.

Charge

Each committee member is asked to review supporting background material and participate in committee discussions. The focus of this review will be on understanding:

- The near- and long-term demand (at a high level) for each hardware architecture in the existing portfolio and how the proposed acquisition will augment or complement the existing hardware portfolio;
- The availability of production software for use by enough of the USQCD collaboration to effectively utilize the capabilities of the proposed acquisition;
- The ability of the proposed acquisition, along with the existing hardware portfolio, to meet the established time-based performance goals for the computing project;
- The capability of the project team to effectively support the computing hardware in the proposed acquisition, in terms of 1) meeting system uptime target goals; and 2) supporting the user community in the use of the newly acquired hardware;
- The alignment of the computing hardware in the existing portfolio and new acquisition with vendor technology roadmaps; and with the technology roadmaps of leadership-class facilities at which USQCD collaboration members run scientific software codes.

Membership

The review committee will comprise members of the LQCD-ext II project and USQCD Collaboration with an appropriate mix of relevant technical and scientific expertise to effectively evaluate the merits of the proposed acquisition plan. The Chair of the committee will be Rob Kennedy. The membership of the FY16 review committee is as follows:

NAME	PROJECT ROLE	AFFILIATION	EMAIL
Gerard Bernabeu	Co-Site Manager (observer)	FNAL	gerard1@fnal.gov
Carleton Detar	Collaboration Representative	U. of Utah	detar@physics.utah.edu
Steve Gottlieb	Collaboration Representative	Indiana U.	sg@indiana.edu
Don Holmgren	Site Architect	FNAL	djholm@fnal.gov
Balint Joo	Collaboration Representative	JLab	bjoo@jlab.org
Rob Kennedy, Chair	Associate Project Manager	FNAL	kennedy@fnal.gov
Shigeki Misawa	Site Architect	BNL	misawa@bnl.gov
Bob Mawhinney	Site Architect	Columbia / BNL	rdm@physics.columbia.edu
James Osborn	Collaboration Representative	ANL	osborn@alcf.anl.gov
Amitoj Singh	Site Manager	FNAL	amitoj@fnal.gov

Deliverables

- A brief, written report summarizing the review committee's analysis of the proposed acquisition plan and assessment of how effectively the proposed plan will meet the computing needs of the scientific program.
- Recommendation(s) to the Project Manager on how best to proceed.

Timeline

The review committee should complete its analysis and provide a written report with recommendations to Bill Boroski, LQCD-ext II Project Manager, no later than May 31, 2016.

Supporting Documentation

The following documentation will be provided to the review committee.

- LQCD-ext II Acquisition Strategy
- LQCD-ext II FY16 Acquisition Plan
- Performance Goals and Milestones for the LQCD-ext II Computing Project
- Anticipated Computing Needs of the Scientific Program (2016-2021)

Requests for additional information should be conveyed through the committee chair to the Project Manager.