

**FY17 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 0.9

May 12, 2017

PREPARED BY:

The FY17 Acquisition Review Committee

*Rob Kennedy, Gerard Bernabeu Altayo, Carleton DeTar, Steve Gottlieb, Don Holmgren, Chulwoo Jung,
James Osborn, Amitoj Singh, Chip Watson, Frank Winter*

FY15 LQCD-ext II User Survey Report Change Log

Version	Description	Date
0.5	Initial draft for review by Committee and Project Manager	5/8/2017
0.9	Revised draft ready for public consumption	5/12/2017
1.0	Final report	<i>Planned: June 2017</i>

Table of Contents

1	Introduction and Background	3
2	Deliverable: Recommendation to the Project Manager	3
3	Deliverable: USQCD-Specific Software Benchmarks	3
4	Deliverable: Addressing Charge Elements	3
5	Suggestion for Future Acquisition Reviews	6
6	Appendix: Charge to the FY17 Acquisition Review Committee	7

1 Introduction and Background

The purpose of the FY17 Acquisition Review Committee is to review and consider the proposed computing hardware acquisition plan for FY17 at BNL 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. This committee was not to reproduce the procurement process at BNL in recommending a best value alternative. Essentially all vendor information considered in this review was public knowledge.

2 Deliverable: Recommendation to the Project Manager

The following recommendation is made by the FY17 Acquisition Review Committee:

The preliminary recommendation of the FY17 Acquisition Review Committee is to proceed with the FY17 LQCD BNL acquisition plan, given the information that we have at this time (5/12/2017).

3 Deliverable: USQCD-Specific Software Benchmarks

A set of USQCD-specific software benchmarks has been developed to evaluate the performance of candidate computing systems proposed by vendors in response to the FY17 acquisition Request for Proposals (RFPs). Bob Mawhinney is compiling the formal list of USQCD-specific software benchmarks that will be used in the BNL FY17 acquisition.

- See Bob Mawhinney's talk at the April 2017 USQCD All-Hands Meeting:
 - <http://www.usqcd.org/meetings/allHands2017/slides/Mawhinney.pdf>
- These benchmarks are also discussed in the FY17 Alternatives Analysis document, in preparation.
- At the time of this writing, the following are proposed to be used to provide a somewhat "portfolio" view of the performance:
 - DWF using Grid, 24^4 local volume, run on 1 and 16 nodes
 - MILC code with optimizations, 32^4 local volume, run on 1 and 16 nodes
 - MILC code with generic C code, 32^4 local volume, run on 1 and 16 nodes

4 Deliverable: Addressing 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:
 - Anna Hasenfratz and Aida El-Khadra, of the USQCD SPC, state that in the 2017 allocation year proposals, the architecture demand was as follows:
 - CPUs are over-requested by a factor of 2.49

- GPUs are over-requested by a factor of 0.98
 - KNLs are over-requested by a factor of 2.19
- Long-term Demand: In the long-term we see:
 - Continued demand for CPU (whether conventional or MIC) technologies
 - A significant fraction of Conventional CPU demand metamorphosing into MIC CPU demand
 - Continued demand for GPU technology
 - Technologies that appear likely to be the most cost-effective for LQCD going forward are:
 - MIC (example today is Intel KNL)
 - GPU (example today is NVIDIA Pascal)
 - CPU may prove cost-effective eventually, but only if certain features like high-bandwidth memory are integrated into the CPU models.
- How will proposed acquisition augment or complement existing hardware portfolio:
 - CPU systems meet the USQCD demand, but do not appear to be sufficiently cost-performant to meet LQCD computing project goals.
 - KNL systems meet the USQCD demand, and are sufficiently cost-performant to meet LQCD computing project goals. We appear to be making progress resolving the operational difficulties seen on KNL systems. For example, Christoph Lehner has successfully run 64-node jobs on JLab's 16p cluster. An upcoming round of BIOS upgrades and chip screening/replacement are expected to improve reliability.
 - The demand for GPU systems appears to already be met; they are less over-requested than CPUs and KNL. GPUs are sufficiently cost-performant to meet LQCD computing project goals provided enough of the software portfolio can run on GPUs.
 - The procedure for summarizing the proposal demands does not take into account that some proposals can be served by CPU/MIC or GPU, so while the demand for GPUs relative to their supply is less than other architectures, they may not be as under-requested as the reported values indicate.

Element 2: Alternate computing architectures that may better meet USCQD needs, considering compatibility with the existing hardware portfolio and infrastructure, as well as the existing software portfolio;

- IBM OpenPower
- ARM ARM64
- The addition of FPGAs to the mainstream Intel architecture may be of interest in the future.
 - We do not see USQCD codes being as well positioned, however, to take advantage of this like some other areas of study though.

Element 3: The availability of production software for use by enough of the USQCD collaboration to effectively utilize the capabilities of the proposed acquisition;

- CPUs: essentially all USQCD software can run on CPUs, provided it is not optimized to some other architecture, though it may not necessarily run at peak performance.
- KNL: most USQCD software can run on KNL, provided it is not optimized to some other architecture. This may require running in a backward-compatibility memory mode, but that can be arranged with the local site administrators for specific jobs.
- GPU: some USQCD software can run on GPUs provided it is optimized for that architecture. Large-memory algorithms however are not as effective on GPUs due to limited total memory per host in the packaging that we could afford.

Element 4: The ability of the proposed acquisition, along with the existing hardware portfolio, to meet the established time-based performance goals for the computing project;

- Acquisition of advanced architectures (MIC, GPU) can meet the time-based performance goals.
- An acquisition of newer CPUs from Intel or AMD alone however will probably not be able to meet these goals due to cost-performance and also possibly due to late delivery.

Element 5: 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;

- While new silicon has historically posed additional risks, we now have some experience on the architectures being considered.
 - There are questions currently about our ability to support very large memory jobs: on GPU systems due to bandwidth, on KNL systems due to operational robustness at large memory footprint scale (which appears to be improving), and on CPUs systems due to cost-performance.
- CPU systems: this is an established architecture. We consider this to have the lowest risk of not meeting uptime goals and the most supportable hardware for the user community. Intel Skylake may arrive late, which could impact its first year delivered computing.
- KNL systems: this has known issues that are being addressed: for example, Fast Memory fragmentation in cache mode, hangs when rebooting to change memory modes. This represents only about a 5-10% performance hit at this time, but it requires additional administrative effort to address these issues. There is no guarantee that the vendor will resolve these issues.
- GPU systems: These have some added risk compared to CPUs, based on experience with the BNL-IC system (NVidia K80s) deployment and Pi0g (NVidia K40s) mid-life, due to packaging and heat management.

Element 6: 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 convergence between conventional Xeon and Xeon Phi line over time.
 - KNL is used by NERSC’s CORI Phase II, Argonne LCF Theta.
 - KNH (Knight’s Hill, successor to KNL) will be used by Argonne LCF Aurora.
 - KNL is in the field, and we have significant production experience with it now.
- NVIDIA Pascal GPUs are also in vendor and LCF roadmaps, and are available now.
 - 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 GPUs.
 - The next generation in the Tesla series, the NVIDIA Volta, has just been announced. It is advertised to deliver 5X the performance of the Pascal. NVIDIA says it will be available later in the year.
- Conventional CPU architectures are not on the LCF roadmaps, but clearly will continue to be evolved by the vendors. Intel Broadwell may become cost-performant if there are “fire sales” as the new Intel Skylake architecture is rolled out. The new Intel Skylake architecture itself is of interest due to its supporting the AVX-512 instruction set. The new AMD Ryzen CPUs appear on paper to be a worthy competitor to the Intel line.

Element 7: Which USQCD-specific software benchmarks should be used in making the best-value assessment during the procurement process.

- See Bob Mawhinney’s talk at the April 2017 USQCD All-Hands Meeting:
 - <http://www.usqcd.org/meetings/allHands2017/slides/Mawhinney.pdf>
- These benchmarks are also discussed in the FY17 Alternatives Analysis document, in preparation.
- At the time of this writing, the following are proposed to be used to provide a somewhat “portfolio” view of the performance:
 - DWF using Grid, 24⁴ local volume, run on 1 and 16 nodes
 - MILC code with optimizations, 32⁴ local volume, run on 1 and 16 nodes
 - MILC code with generic C code, 32⁴ local volume, run on 1 and 16 nodes

5 Suggestion for Future Acquisition Reviews

Suggestion: It would be useful to document the software portfolio by creating a list of USQCD production jobs with their performance characteristics and resource requirements. This will require effort from both users and the site managers to accomplish.

6 Appendix: Charge to the FY17 Acquisition Review Committee

LQCD-ext II FY17 Acquisition Review Committee Charge

February 1, 2017

Revised February 20, 2017

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 hardware portfolio is used by USQCD in support of its scientific program.

In fiscal year 2017, the project will execute two distinct hardware acquisitions.

1. The first acquisition has already occurred: \$410K in project hardware funds have been used to purchase an extension of the Knights Landing cluster implemented at JLab last year, as outlined and in accordance with the FY16 Hardware Acquisition Plan. The purchase order was executed in October and the new hardware has been received, installed, and is in production use.
2. The second acquisition will occur in calendar year 2017, wherein up to \$750K in project hardware funds will be used to purchase and deploy new computing hardware at Brookhaven National Laboratory (BNL). Following standard project practices, preparation of the FY17 Acquisition Plan is underway and is being led by Bob Mawhinney, one of the LQCD-ext II co-Site Architects for BNL, with input from the rest of the project team.

The purpose of this committee is to review and provide input into the FY17 computing hardware acquisition planning process to ensure strong alignment of the hardware procurement with the anticipated computing needs of the USQCD scientific program. The intent is to help ensure that the project is making the most effective use of computing hardware funds to support and advance the scientific program.

Charge

The Acquisition Review Committee is asked to review the FY17 Acquisition Plan, along with supporting information and materials, and prepare a brief report summarizing the committee's assessment of the plan and proposed hardware acquisition. Each committee member is asked to review supporting materials, provide input, actively participate in committee discussions, and where possible, provide USQCD-specific code samples that will help the project benchmark the performance of candidate hardware against USQCD needs. The committee is asked to consider:

- 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;
- Alternate computing architectures that may better meet USCQD needs, considering compatibility with the existing hardware portfolio and infrastructure, as well as the existing software 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;
- Which USQCD-specific software benchmarks should be used in making the best-value assessment during the procurement process.

The committee may participate in activities leading up to the evaluation of vendor responses to the RFP, but will not participate in the vendor proposal evaluation process, nor the selection of the best-value proposal. The LQCD-ext II Computing Project adheres to the procurement policies and procedures of the host laboratory for the acquisition, which impose constraints on who can serve on the procurement evaluation committee (typically employees of the host institution).

Deliverables

- A set of USQCD-specific software benchmarks that can be used to evaluate the performance of candidate computing systems proposed by vendors in response to the FY17 acquisition Request for Proposals (RFPs).
- 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 with the hardware acquisition.

Timeline

- The review committee should assemble a set of appropriate benchmarks that can be included in the acquisition RFP by February 24, 2017. This will provide the project with sufficient time to include the benchmarks in the RFP documentation; the RFP is scheduled to be released on March 1, 2017.
- The review committee should complete its full analysis and provide a final written report with recommendations to Bill Boroski, LQCD-ext II Project Manager, no later than March 17, 2017.

Membership

The review committee comprises 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. In addition, the committee includes a member of the USQCD Scientific Program Committee, so the interests, needs and input of that committee are represented and factored into the process.

The Chair of the committee will be Rob Kennedy. The membership of the FY17 review committee is as follows:

NAME	PROJECT ROLE	AFFILIATION	EMAIL
Gerard Bernabeu	Site Manager	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	HPC Consultant	FNAL	djholm@fnal.gov
Chulwoo Jung	Collaboration Representative	BNL	chulwoo@bnl.gov
Rob Kennedy, Chair	Associate Project Manager	FNAL	kennedy@fnal.gov
James Osborn	Collaboration Representative	ANL	osborn@alcf.anl.gov
Amitoj Singh	Site Architect	FNAL	amitoj@fnal.gov
Chip Watson	Site Architect	JLab	watson@jlab.org
Frank Winter	Collaboration Representative	JLab	fwinter@jlab.org
Alex Zaytsev	Site Architect	BNL	alezayt@bnl.gov

Supporting Documentation

The following documentation will be provided to the review committee as documents and information becomes available.

- LQCD-ext II Acquisition Strategy
- LQCD-ext II FY17 Acquisition Plan
- Performance Goals and Milestones for the LQCD-ext II Computing Project
- Anticipated Computing Needs of the Scientific Program (2017-2021)
- Performance data on USQCD applications running on the actual FY16 production hardware to compare actual performance against early benchmarks.

Requests for additional information should be made to the chairperson of the review committee.