

**Risk Management Plan
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 1.1

March 27, 2014

PREPARED BY:
Robert D. Kennedy, FNAL

CONCURRENCE:



William Boroski
LQCD-ext II Contract Project Manager

April 24, 2015
Date

**LQCD-ext II Risk Management Plan
Change Log**

Revision No.	Description/Pages Affected	Effective Date
0	Adopt this document for LQCD-ext II from LQCD-ext	February 12, 2014
1.0	Medium priority risks are reviewed quarterly. Add “Next Review Date” field to register, process.	March 27, 2014
1.1	Review (no changes), update signature date	April 24, 2015

Table of Contents

1	Introduction.....	3
2	Overview of the Risk Management Plan	3
2.1	Purpose and Objective	3
2.2	Responsibility	4
3	Risk Assessment	4
3.1	Risk Assessment Planning	4
3.2	Execution of Risk Management.....	4
3.3	Re-plan.....	5
4	Risk Identification.....	5
5	Risk Analysis	6
6	Risk Handling	6
6.1	Technology Risk Mitigation	7
6.2	Cost Risk Mitigation	7
6.3	Schedule Risk Mitigation.....	8
6.4	Security Risk Mitigation	8
6.5	Service Risk Mitigation	8
7	Risk Monitoring.....	9
8	List of Acronyms	10

1 Introduction

This document describes the risk management plan for the SC Lattice Quantum Chromodynamics (LQCD) Computing Project Extension II, referred as LQCD-ext II in the rest of the document. This document and the associated Risk Registerⁱ are LQCD-ext II Controlled Documents.

The purpose of LQCD-ext II Computing project is the deployment and operation of a large scale dedicated computing facility capable of delivering an aggregate of over 220 Tflop/s-yrs (1 Tflop/s or one teraflop per second = 10^{12} floating point operations per second) by the end of FY2014 for the study of Quantum Chromodynamics (QCD). This project plays an important role in expanding our understanding of the fundamental forces of nature and the basic building blocks of matter. The computing hardware is housed at Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (FNAL) and Thomas Jefferson National Accelerator Facility (TJNAF), and is operated as a single distributed computing facility, which is available to lattice gauge theorists at national laboratories and universities throughout the United States of America.

LQCD-ext II operates a number of discrete cluster or other computing resources at each of the host laboratories. Since the project designs, procures, and commissions one or more new resources in each year, and retires existing resources after approximately four years of operation, the precise portfolio varies from year to year. In any given year, across the three host laboratories, the project operates multiple conventional Infiniband-based clusters, multiple Infiniband-based clusters accelerated with general-purpose graphics processing units (GP-GPUs), and in some years BlueGene-architecture supercomputer hardware. The project has also recently added a small Xeon Phi cluster.

2 Overview of the Risk Management Plan

2.1 Purpose and Objective

As defined in the LQCD-ext II Project Execution Plan, the Integrated Project Team (IPT) views risk management as an ongoing task that is accomplished using a formalized plan, namely this document, to identify, analyze, mitigate and monitor the risks that arise during the course of the project. LQCD-ext II established its risk management plan during the early stages of the project using the guidelines set forth in Chapter 14 of DOE Publication M 413.3-1ⁱⁱ, Project Management for the Acquisition of Capital Assets. The current revision of the document is based on the guidance provided in the Guide to the Project Management Body of Knowledgeⁱⁱⁱ and the OMB Circular Number A-11 Part 7 Capital Programming Guide^{iv}.

As defined in these references, risk is a measure of the potential of failing to achieve overall project objectives within the defined scope, cost, schedule and technical constraints. The purpose of this document is to describe how LQCD-ext II IPT plans to minimize the project risks and document actions to put in place in a timely and consistent manner in case of an occurrence. The LQCD-ext II risk management strategy is to avoid risk as much as possible by understanding the possible risks associated with the project and devising methodologies for managing them. LQCD-ext II risks can be envisioned from two points of view:

- a. Enclave-based risks associated with the hosting laboratories, including security, privacy and business continuity: LQCD-ext II equipment hosted by each laboratory remains under the

jurisdiction of the hosting laboratory as a part of its General Enclave. Security, privacy and business continuity risk, also called disaster recovery, responsibilities are managed by the hosting laboratory.

- b. Project-based risks associated with the overall project: Project-based risks are the risks associated with the overall project execution.

Enclave-based risk management methods are addressed in the laboratory specific documents and technologies.

2.2 Responsibility

The final responsibility for risk management rests with the Contract Project Manager (CPM), who takes appropriate measures in consultation with the LQCD-ext II Integrated Project Team, the Executive Committee and other project members. Designated Site Managers at each site are responsible for the site specific risks. However, effective risk management is an iterative, multi-step process that requires the continued involvement of all project members.

3 Risk Assessment

Since the goal of the LQCD-ext II project is to extend the envelope of technology, specifically to establish and operate dedicated systems that optimize performance/price ratio for LQCD computing, it is necessary to accept certain levels of risks to achieve the scientific objectives of the project. Eliminating risk entirely is not a viable option. The LQCD-ext II IPT adopted a “risk aversion to a moderate degree” approach. The strategy is to reduce risk to an acceptable level by using the project plan effectively to mitigate risks as they arise. The project uses various control mechanisms to manage residual risks. The risk management process of LQCD-ext II is integrated with the technical plans, the Project Execution Plan and the Work Breakdown Structure for the project.

3.1 Risk Assessment Planning

The technical plan for the project, as documented in the associated acquisition plans, hardware benchmarks, and alternate strategies, outlines the risks associated with the annual investment and their impacts. This planning process includes identification of risks, probability of occurrence, degree of impact, and risk mitigation strategy. A change management process, as outlined in the LQCD-ext II Project Execution Plan, is in place to manage changes to the project that may occur to mitigate realized risks. Identified risks are documented in the LQCD-ext II Risk Register which also contains records of outcomes of the qualitative risk assessment.

3.2 Execution of Risk Management

As the project progresses, the LQCD-ext II IPT evaluates the risks continuously by using project management metrics and tools including:

- Monthly project completion status reports
- Monthly financial status reports
- Monthly technical accomplishment reports
- Change requests and their approvals or rejections

The LQCD-ext II IPT reviews risks and its net risk levels of the project continuously. If a decision is made that the net risk level of a particular item requires that a risk mitigation strategy should be put in place, then any associated changes required by the mitigation is handled according the the Project Change Control procedure described in the LQCD-ext II Project Execution Plan.

3.3 Re-plan

Mitigation plans for new risk assessment results are incorporated into the plans for subsequent years. If necessary, a change control request is also processed. The current list of mitigation plans is in the Risk Register.

4 Risk Identification

The Risk Register workbook contains a tab “Risk Register” with the list of risks identified for the project, their attributes and their risk ratings. Attributes associated with each risk are as follows:

- Risk ID – unique identifier across all risks
- Risk Area – for risk categorization, one of:
 - Cost, Schedule, Technology; Security, Service
- Risk Title – Short description of the risk for reporting
 - No more than 3 lines at current field width.
- Description – Long description of the risk (narrative)
- Probability of Occurrence – one of:
 - High, Medium, Low
- Impact of Occurrence – one of:
 - Severe, Moderate, Low
- Risk Rating – numerical value based on Risk Rating Table
 - Rating Value = Probability Value * Impact Value
- Risk Priority – one of: (based on Risk Rating value)
 - High, Medium, Low
- Risk Status – one of:
 - Active – risk condition has occurred, so the risk has become a real issue.
 - Exists – risk condition has not occurred
 - Retired – risk condition is no longer considered worthwhile tracking
- Creation Date (date)
- Last Review Date (date)
- Next Review Date (date) – driven by review frequency or earlier date chosen by IPT
- Last Change (narrative, should include date)
- Mitigation Strategy (narrative)
- Notes (narrative)

Detailed information regarding each identified risk is recorded in the narrative fields in the Risk Register.

5 Risk Analysis

Each identified risk for the project is analyzed for the probability and impact of occurrence. Individual ratings for probability and impact of occurrence are assigned to each of them based on the values shown in Table 1. The risk rating is then derived by multiplying probability and impact values shown in Table 2. Finally, the risk is assigned a priority based on the risk rating value which will drive the level of planning and frequency of monitoring that will be used for the risk.

Table 1: Risk Probability and Impact Values

Probability	Value	Impact	Value
High	0.75	Severe	0.9
Medium	0.50	Moderate	0.5
Low	0.25	Low	0.1

Table 2: Risk Rating Matrix (with Risk Priority color coding)

Prob \ Impact	Severe	Moderate	Low
High	0.675	0.375	0.075
Medium	0.450	0.250	0.050
Low	0.225	0.125	0.025

Table 3: Risk Prioritization Table

Risk Priority	Rating Low Value	Rating High Value	Risk Planning Level	Risk Plan Location	Risk Review (Minimum) Frequency
1 - High	0.500	1.000	Detailed Risk Plan	Separate Document	At least monthly
2 - Medium	0.150	0.500	Modest Risk Plan	Risk Register	At least quarterly
3 - Low	0.000	0.150	Minimal Risk Plan	Risk Register	At least annually

6 Risk Handling

The primary risk handling strategy for the LQCD-ext II project is to avoid risks where possible by making good project assumptions and then validating those assumptions with the DOE Review Committee and Scientific Program Committee. The LQCD-ext II project also uses various risk minimization tools and techniques, such as:

- System and subsystem prototyping
- Benchmarking using modeling and simulation
- Formal and informal technology assessments
- Quality control and system validation
- Alternative acquisition analysis
- System and subsystem level risk assessments including prioritization
- Continuous monitoring of technical and financial performance measures
- Establishing various surety measures including security and disaster recovery measures

Risks for the project are categorized into one of five major “risk areas”: Technology, Cost and Schedule (following DOE guidelines), as well as Security (computer security and privacy issues), and Service (business/service continuity and disaster recovery issues). The risk mitigation strategies for each risk categories are developed to minimize and mitigate the risks involved with most project deliverables. General mitigation strategies for each risk area are described below.

6.1 Technology Risk Mitigation

The major technical concern for the LQCD-ext II project is the annual delivery of computing capabilities, expressed in Tflop/s-yrs. Since this is related in part to the cost of the new systems, and the schedule for their delivery, the risk involved is of low to moderate probability and with moderate impact. In any given year, the computing capacity of the new system commissioned in that year will not exceed 30% of the total computing capacity available to the project. Further, each new system is planned to be operated for at most the last 3 months of a given fiscal year, except possibly for the last year of the project. Consequently LQCD-ext II can reliably predict prior to the beginning of any fiscal year the Tflop/s-yrs that will be delivered in the fiscal year. This allows for detailed planning, by the Scientific Program Committee, of allocations to scientists for access to these computing resources. It is also possible to track and benchmark new products available in the market.

Possible schedule overrun is also a risk associated with technology. The risk of schedule overrun by the LQCD-ext II project is of low to moderate probability and of moderate impact. The schedule estimates are based on the promised release dates (“roadmaps”) for hardware components as given by the manufacturers, and the delivery dates given by the third-party vendors and integrators with whom the LQCD-ext II project subcontracts for the hardware purchases. Since the LQCD-ext II project must rely on state-of-the-art technologies to deliver highest possible computing power within the project budget, it is often necessary to wait for the most advanced technologies, for example, processor and switching technologies, promised by the manufacturer. However, if the manufacturer fails to make good on the promised dates, the schedule may slip, or the project may have to procure the existing technology at lower performance.

6.2 Cost Risk Mitigation

Because LQCD-ext II funding is directly associated with the Congressional release of funds, there may be a delay in the availability of moneys for major procurements. To mitigate this risk, all major LQCD-ext II procurements are scheduled after the end of first quarter of each fiscal year. The risk of cost overrun or exceeding the available fund by the LQCD-ext II project is of low probability and of low impact. The cost estimates are based in part on previous procurements for the SciDAC prototype systems, procurements during the LQCD-ext II project, and the actual costs of labor for deploying and operating the LQCD-ext II project systems. Together, these firmly establish the historical performance and price trends for COTS-based parallel computing systems for LQCD-ext II calculations. Because of the build-to-cost nature of the project, LQCD-ext II has minimal risk for completing over budget. Hardware cost variances from the estimates described above will result in adjustments to the sizes of the computing systems developed each year. That is, higher than anticipated hardware costs will result in the procurement of a smaller cluster in a given year, or a cluster of different composition (for example, selection of high performance network and/or processor). Labor cost variances, for example, the need to change the amount of user support, will result in adjustments of the division between subsequent equipment and labor budgets. The performance risks associated with computing and network system are estimated to be low due to the successful R&D during the SciDAC project, and the use of COTS hardware wherever possible. Further, the use of conservative extrapolations from

historical LQCD computing performance trends mitigates the risk of delivering less capable computing systems than planned.

Staffing issues may also affect the project cost. Since only a small number of technical staff members are associated with the LQCD-ext II project, there is a low probability of risk associated with the loss of key project members. However, the impact of the loss of key personnel can be high in terms of full release of new computing systems to the scientific community and annual technical delivery. To mitigate this risk, as much as practical, LQCD-ext II staff members at two or more of the host sites participate in the prototyping, planning, and execution of each major system acquisition. Cross-training of system administration duties is encouraged whenever possible. This ensures that LQCD-ext II maintains project expertise in at least several individuals.

6.3 Schedule Risk Mitigation

Schedule risk for the LQCD-ext II project is tied to the technology and the cost risks. If technology does not keep up with the technology roadmap or if the cost becomes higher than projected, the primary scientific delivery schedule will be affected. The general risk mitigation strategy in this area consists of optimizing the annual procurements, both from cost and schedule point of view, to deliver the most TFlop/s-yrs.

6.4 Security Risk Mitigation

Security management of the computing facility hosted by each laboratory is administered by the physical and cyber-security infrastructure established by that laboratory. The LQCD-ext II Security Plan for a particular set of computing equipment within an enclave is updated and approved whenever any new equipment is added to the existing enclave. Each laboratory must also keep its Certification and Accreditation (C&A) documents up-to-date. The hosting laboratory also performs required scans and other monitoring and assessments. Since LQCD-ext II computing facility is special purpose equipment dedicated to LQCD simulation, there is very little external access to the equipment. No private, personal or otherwise, information may be retained on the LQCD-ext II computer facilities.

6.5 Service Risk Mitigation

Since delivering technical results to the USQCD user community is the most critical objective of this investment, the LQCD-ext II project has considered disaster recovery planning from the beginning of the project. The LQCD-ext II project takes advantage of the institutional disaster recovery plans for the computing centers at each of its Laboratory sites. These plans are reviewed periodically. The most valuable data products produced by the project are the vacuum gauge configuration data files, which may require in aggregate many Tflop/s-yrs of computing. These files are stored redundantly at multiple locations, including two or more of FNAL, TJNAF, BNL, NERSC, TACC, ORNL, the ALCF, and LLNL. The principal investigator for each computational project executed on the LQCD-EXT II systems is responsible for safeguarding the data products produced by his or her scientific project. By standard government policy, the equipment at each facility will not be insured against disasters, though the standard safety protections provided by each laboratory assure as much as possible the protection of the equipment. The distributed nature of the meta-facility partially mitigates the risk of natural disasters, allowing for critical scientific calculations to be moved from one host site to another in the event of a sustained outage.

7 Risk Monitoring

The Risk Register is reviewed and updated whenever a new risk is identified, a risk is determined to no longer affect the project, or a change in the environment is detected that may affect existing risks or risk plans. In steady state, absent such changes, a graded approach will be used for risk monitoring based on the risk priorities summarized in Table 3.

Risks assigned a “1 – High” priority will receive the most attention in risk planning. They will be reviewed at least monthly and whenever the IPT believes there is a need to review the risk, as tracked in the Next Review Date field. They will have a detailed risk plan developed and maintained in a separate document outside of the Risk Register (a document reference is stored in the Risk Register).

Risks assigned a “2 – Medium” priority will receive moderate attention in risk planning. They will be reviewed at least quarterly, and whenever the IPT believes there is a need to review the risk, as tracked in the Next Review Date field. They will have modest risk plans which will be captured in the narrative fields of the Risk Register.

Risks assigned a “3- Low” priority will receive the least attention in risk planning. They will be reviewed annually, and whenever the IPT believes there is a need to review the risk, as tracked in the Next Review Date field. They will have minimal risk plans, often just noting that the risk is accepted, which will be captured in the narrative fields of the Risk Register.

LQCD-ext II Project Office reviews the Risk Register monthly for entries that may be affected by changes in the environment or which have an upcoming Next Review Date to prepare the IPT for a risk item review. The LQCD-ext II IPT will also take advantage of the annual DOE Progress Review to review the long-term risk management plans with the reviewers. The LQCD-ext II Risk Management Plan is updated annually.

8 List of Acronyms

Acronym	Definition
ALCF	Argonne Leadership Computing Facility
AMD	Advanced Micro Devices (a processor company)
BNL	Brookhaven National Laboratory
C&A	Certification and Accreditation (computer security)
CCB	Change Control Board
COTS	Commercial Off-The-Shelf
CPU	Central Processing Unit
DOE	Department of Energy
FNAL	Fermi National Accelerator Laboratory (a.k.a. Fermilab)
IPT	Integrated Project Team
LLNL	Lawrence Livermore National Laboratory
LQCD	Lattice Quantum Chromodynamics
NCSA	National Center for Supercomputing Applications
NERSC	National Energy Research Scientific Computing Center
ORNL	Oak Ridge National Laboratory
QCD	Quantum Chromodynamics
QCDOC	QCD On a Chip (BNL Supercomputer)
SciDAC	Scientific Discovery through Advanced Computing
TACC	Texas Advanced Computing Center
TFlop/s	Teraflops per second. 1 teraflop = 10 ¹² floating point operations
TFlop/s-yr	Computing delivered by 1 TFlop/s sustained for one year
TJNAF	Thomas Jefferson National Accelerator Facility (a.k.a. JLab)

ⁱ Risk Register for the LQCD-ext II project

ⁱⁱ DOE G 413.3-7 Risk Management Guide (9-16-08)

ⁱⁱⁱ A Guide to the Project Management Body of Knowledge (PMBOK® Third Edition), Project Management Institute

^{iv} OMB Circular Number A-11 Part 7 Capital Programming Guide V2.0 (2006) Appendix 5