

**Risk Management Plan
for the
SC Lattice QCD Computing Project (LQCD)**

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

July 17, 2007

Revision 2.0

PREPARED BY:
Bakul Banerjee and
Don Holmgren, FNAL

CONCURRENCE:



William Boroski
LQCD Contract Project Manager

7-17-2007
Date

Lattice QCD Computing Project (LQCD) Risk Management Plan

Change Log

| Revision No. | Description/Pages Affected | Effective Date |
|--------------|--|----------------|
| Revision 0.0 | Risk assessment / risk management plan for BY06 OMB Exhibit 300 | July 1, 2004 |
| Revision 0.1 | Risk management / Risk assessment incorporated into Project Execution Plan | May 1, 2005 |
| Revision 0.2 | Additional risks added (ID 19, 20) | June 1, 2005 |
| Revision 0.3 | Additional risk added (ID 21) | August 8, 2005 |
| Revision 0.4 | Technical obsolescence risk modified (ID 4) | Nov 1, 2005 |
| Revision 0.5 | Risk management plan broken out of Project Execution Plan | Dec 1, 2005 |
| Revision 0.6 | Feasibility risk modified (ID 5) | Dec 20, 2005 |
| Revision 0.7 | Update for BY08 OMB Exhibit 300 submission | April 6, 2006 |
| Revision 0.8 | Technical obsolescence risk modified (ID 4) | May 25, 2006 |
| Revision 0.9 | Update of all risk status items / pp 8 – 18 | Sept 25, 2006 |
| Revision 1.0 | Final FY06 Risk Management Plan | Sept 25, 2006 |
| Revision 2.0 | Final FY07 Risk Management Plan – entire document | July 17, 2007 |

Table of Contents

1. INTRODUCTION 1

2. OVERVIEW OF THE RISK MANAGEMENT PLAN 1

 2.1 PURPOSE AND OBJECTIVE1

 2.2 RESPONSIBILITY.....2

 2.3 RISK ASSESSMENT2

3. RISK IDENTIFICATION 4

4. RISK ANALYSIS..... 5

5. RISK HANDLING 5

 5.1 GENERAL RISK MITIGATION5

 5.2 DISASTER RECOVERY6

 5.3 RISK MINIMIZATION AND MITIGATION TOOLS AND TECHNIQUES6

6. IDENTIFIED RISKS, MITIGATION STRATEGIES, AND STATUS 6

7. RISK MONITORING..... 18

8. REFERENCES 18

APPENDIX A: SUMMARY OF RISK RATINGS BY CATEGORY 19

APPENDIX B: GLOSSARY 20

1. INTRODUCTION

The purpose of this document is to describe the risk management plan associated with the SC Lattice Quantum Chromodynamics (LQCD) Computing project, referred as LQCD in the rest of the document, and the annual risk management status updates. This document and the associated Risk register are LQCD Controlled Documents.

The purpose of LQCD Computing project is the deployment and operation of a large scale dedicated computing facility capable of sustaining over seventeen (17) Tflop/s (teraflop per second, where 1 teraflop = 10^{12} floating point operations) 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. The project started in October 1, 2006 and will be completed in FY 2009. The total project cost is \$9.2M.

LQCD operates a number of systems which were already in existence at the beginning of the investment: (1) the QCDOC purpose-built supercomputer at BNL, (2) two commodity clusters at FNAL constructed as large scale production prototypes during the DOE's SciDAC Lattice QCD Computing Project, and (3) two additional commodity clusters at TJNAF, also from the SciDAC Project. Two additional SciDAC clusters, one at FNAL and one at TJNAF, reached end-of-life during the first year of this investment. LQCD also deploys new systems. During each year of the four-year project, LQCD will design, procure, and commission a new system at either TJNAF or FNAL. In the first year, FY06, LQCD commissioned a small system at TJNAF, and a large system at FNAL. IN FY07, LQCD deployed a large system at TJNAF. During BY08 and BY09, a large combined cluster will be deployed at FNAL using hardware fund allocations for two years. These newly deployed systems are composed of commercial off-the-shelf (COTS) hardware interconnected with high performance networks such as Infiniband.

2. OVERVIEW OF THE RISK MANAGEMENT PLAN

2.1 PURPOSE AND OBJECTIVE

As defined in the LQCD 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 established its risk management plan during the early stages of the project using the guidelines set forth in Chapter 14 of DOE Publication M 413.3-1, Project Management for the Acquisition of Capital Assets. The current revision of the document is based on the guidance provided in *A Guide to the Project Management Body of Knowledge (PMBOK Guide, Third Edition)* and OMB Circular Number A-11 Part 7 Capital Programming Guide V2.0 (2006) Appendix 5.

As defined in above 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 plan is to describe how the LQCD IPT plans to minimize project risks and document actions that may be put into place in a timely and consistent manner in case of an occurrence. The LQCD 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 uses key procedures proven to be an effective strategy in the risk management of scientific projects: planning, assessment, handling and monitoring.

2.2 RESPONSIBILITY

The final responsibility for risk management rests with the Contractor Project Manager (CPM). The CPM is assisted in the identification and management of project risk by the IPT, LQCD Executive Committee, and other project members. Designated Site Managers at each of the host laboratories are responsible for identifying and managing site-specific risks, in consultation with the CPM as appropriate.

2.3 RISK ASSESSMENT

Since the goal of the LQCD project is to extend the envelope of the technology, specifically to establish and operate dedicated systems that optimize the performance/price ratio for LQCD computing, it is necessary to accept certain levels of risks to achieve the scientific objectives of the project. Therefore, the LQCD IPT has adopted a “risk aversion to a moderate degree” approach. The strategy is to reduce risk to an acceptable level by using project planning to effectively manage and mitigate risks as they arise. The project uses various control mechanisms to manage residual risks. The risk management process is integrated with LQCD technical plans, the Project Execution Plan, and the Work Breakdown Structure.

- Risk assessment planning: The technical plan for the project, as documented in hardware acquisition plans, hardware benchmarking results, and alternatives analysis, outlines the risks associated with the investment and their respective potential impacts. The planning process includes identification of risks, probability of occurrence, level of impact, and risk mitigation strategy. A change management process, as outlined in the LQCD Project Execution Plan, is in place to manage changes to the project that may need to occur to mitigate realized risks. Identified risks are documented in the LQCD Risk Register, which also records outcomes of the qualitative risk assessment. Details of LQCD project risks are given in the risk identification section of this document.
- Execution of risk assessment: As the project progresses, the LQCD IPT continuously evaluates risks by using project management metrics and tools, including:
 - Monthly project completion reports
 - Monthly financial status reports
 - Monthly technical accomplishment reports
 - Change requests and their approvals or rejections

In consultation with the IPT, the CPM reviews risks and the net risk level of the project regularly. If a decision is made that the net risk level of a particular item requires that a

risk mitigation strategy be put into place, then the required change is evaluated against the change control threshold table in the LQCD Project Execution Plan. If the change level is low, then the site managers and the project manager executes necessary changes needed to mitigate the risks. If the risk level is high, and the change needed to implement the risk mitigation strategy requires the approval of the LQCD Change Control Board (CCB), then a Change Control Request is issued. Any significant changes needed to mitigate risks are approved by the CCB.

- Review of assessed risks: During the third quarter of each fiscal year, an external DOE Progress Review committee assesses the LQCD project. The LQCD IPT takes advantage of this assessment process to review long-term risk management plans with the reviewers.
- 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 list of mitigation plans is given in the risk identification section of this document.

The LQCD project has three interrelated, major risk areas, namely, cost, schedule and technical deliverables. The risks and mitigation strategies are described below.

Cost: The risk of cost overrun by the LQCD project is of low probability and low impact. Project cost estimates are based in part on previous procurements for the SciDAC prototype systems, procurements during the LQCD project, and the actual costs of labor for deploying and operating the SciDAC and LQCD computing systems. Together, these firmly establish the historical performance and price trends for COTS-based parallel computing systems for LQCD calculations. Because of the build-to-cost nature of the project, LQCD 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 systems are estimated to be low due to successful R&D efforts 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.

Schedule: The risk of schedule overrun by the LQCD project is of low to moderate probability and of moderate impact. Schedule estimates are based on the promised release dates (“roadmaps”) for hardware components as given by component manufacturers, and the delivery dates given by the third-party vendors and integrators with whom the LQCD project subcontracts for hardware purchases. Since the LQCD project must rely on state-of-the-art technologies to deliver the 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 a manufacturer fails to make good on promised release dates, the schedule may slip, or the project may have to procure the existing technology at lower performance. Another related schedule risk is a delay in the availability of project

funding. Because LQCD funding is directly associated with the Congressional release of funds, there may be a delay in the availability of funds for major procurements. To mitigate this risk, all major LQCD procurements are scheduled after the end of the first quarter of each fiscal year.

Technology: The major technical concern for the LQCD 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 of moderate impact. In any given year, the computing capacity of the new system commissioned in that year is not anticipated to 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 can reliably predict, prior to the beginning of any fiscal year, the Tflop/s-yrs that will be delivered in that fiscal year. This allows for detailed planning, by the LQCD 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.

Other areas: Other significant areas of risk identified for the project are:

- Business
- Data
- Organizational and change management
- Privacy
- Project management
- Security

These are described and discussed in detail in Section 6.

3. RISK IDENTIFICATION

The LQCD project has initiated a Risk Register using Microsoft Excel. The risk register workbook contains multiple worksheets that include the list of risks identified for the project, with their respective attributes and risk ratings. Attributes associated with each risk are as follows:

1. Initial Risk ID
2. Risk Area
3. Description
4. Initial date of identification
5. Last update
6. Probability of occurrence of the risk (latest & historical)
7. Impact of occurrence of the risk (latest & historical)
8. Risk rating (probability * impact)
9. Status of monitoring

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

4. 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 risk item. Numerical values assigned to each probability and impact category are shown in Table 1 shows the numerical values assigned to the various probability and impact categories. Table 2 shows the resulting risk ratings, which are derived by multiplying the probability and impact values.

Table 1: Values of risk probability and impacts

| 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 ratings

| | Severe | Moderate | Low |
|---------------|---------------|-----------------|------------|
| High | 0.675 | 0.375 | 0.075 |
| Medium | 0.45 | 0.25 | 0.05 |
| Low | 0.125 | 0.125 | 0.025 |

5. RISK HANDLING

The primary risk handling strategy for the LQCD project is to avoid, or at least minimize, risks by making sound procurement plans and decisions based on the best possible assumptions, and by validating those assumptions using input from the external DOE Progress Review Committee and the internal LQCD Scientific Program Committee. Recognizing that full risk avoidance is not possible, we strive to minimize the impact of potential risks by identifying mitigation actions on a case-by-case basis.

5.1 GENERAL RISK MITIGATION

The risk mitigation strategies for cost and schedule can help mitigate the risks associated with most project deliverables. For example, with regard to manpower resources, there is a low probability of risk associated with the loss of key project members. Since only a small number of technical staff members are associated with the LQCD project, the impact of such a personnel loss could potentially be high in terms of the timely release of new computing systems to the scientific community, and the achievement of annual technical performance goals and milestones. To mitigate this risk, LQCD staff members at two or more of the host sites participate in the prototyping and planning of each major system acquisition. In addition, during system deployment, we strive to maintain frequent dialogue and communication between staff members at the various host institutions. This helps ensure that LQCD maintains project expertise in at least several individuals.

5.2 DISASTER RECOVERY

Since delivering technical results delivered to the USQCD user community is the most critical objective of this investment, the LQCD project has taken into account disaster recovery planning since project inception. The LQCD project takes advantage of the institutional disaster recovery plans for the computing centers at each of its national laboratory host sites.

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 the following sites: FNAL, TJNAF, NERSC, or NCSA. The principal investigator for each computational project executed on the LQCD 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, by allowing for critical scientific calculations to be stored at multiple sites, or to be moved from one host site to another in the event of a sustained outage.

5.3 RISK MINIMIZATION AND MITIGATION TOOLS AND TECHNIQUES

The LQCD project 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.

6. IDENTIFIED RISKS, MITIGATION STRATEGIES, AND STATUS

In this section, detailed descriptions of each risk, their status, and the individual mitigation actions are described. Details of risk attributes are recorded in the Risk Register. If there is a change in the probability and impact of occurrence of the risk, the status and ratings of the risk is updated and the mitigation strategy is revised accordingly. Note that these risks do not appear in any prioritized order.

Risk #1: Delays in New System Deployment

Date Identified: 7/1/2004

Area of Risk: Schedule

Description: The schedule for achieving LQCD investment milestones might slip for the following reasons: A) Vendors may take longer than anticipated to bring new processors, memory systems, and/or interconnect systems to market; B) It may take longer than expected to bring new systems on line for production use; or C) Funding may be lower than anticipated.

Probability of Occurrence: Medium

Impact of Occurrence: Moderate

Risk Rating: 0.25

Mitigation Strategy: A research and development program on the design and implementation of cost effective parallel computing hardware and software for LQCD computations was begun in 2001 as part of the DOE SciDAC Lattice QCD Computing Project. This five-year project was renewed as part of the DOE SciDAC-2 program and will thus continue throughout the lifetime of this project. Experienced professional staff are following the commodity market carefully and gaining insight by evaluating prototype hardware. Project members frequently meet with vendors under non-disclosure agreements and are briefed on roadmaps for components such as processors, chipsets, motherboards, network interface cards and switches. In addition, working closely with manufacturers and system integrators has resulted in the testing of prerelease components. This has both allowed the manufacturers to be informed of deficiencies in their products, and the LQCD team to determine whether some new capability will actually provide any advantage in future systems. As with any investment, a successful implementation of the schedule assumes the approved Budget Authority profile.

Status: No change.

Annual reviews are completed by June 30 of each year of the investment to validate planned modifications to the project baseline in response to schedule slips due to the various areas listed. The DOE held a project cost and schedule review on May 24-25, 2005; the reviewers supported and recommended the project's strategy and schedule for FY06. Subsequent annual DOE project progress reviews were held on May 25, 2006 and May 14-15, 2007; the reviewers again supported and recommended the project's strategy and procurement schedule for the FY07, FY08 and FY09 deployments.

Risk #2: Inaccurate Cost Projections

Date Identified: 7/1/2004

Area of Risk: Cost

Description: Although cost projections for the current budget year are reasonably reliable, projections for subsequent years become progressively uncertain. This may lead to cost overruns.

Probability of Occurrence: Medium

Impact of Occurrence: Moderate

Risk Rating: 0.25

Mitigation Strategy: Market information will be gathered and prototypes built throughout the lifetime of the project. Open procurements of commodity components will provide competitive prices. All hardware is modular in nature, so if prices exceed expectations in any given year, it is possible to deploy smaller machines.

Status: No change.

In each year of the investment we review the cost and performance projections for the next year and present an acquisition plan to an external review panel for that coming year's purchases. As noted before, the process of gathering market information and results were reviewed during the DOE Project Progress Review held in May 2006.

Risk #3: Increases in Life-Cycle Costs**Date Identified:** 7/1/2004**Area of Risk:** Cost**Description:** Unexpected increases in life-cycle costs arise after systems are acquired, leading to cost overruns.**Probability of Occurrence:** Low**Impact of Occurrence:** Moderate**Risk Rating:** 0.125**Mitigation Strategy:** Hardware maintenance costs are included in the procurement of components for each new system acquired. Operations costs are well understood, as they are based on years of similar operating experience. Each of the three host institutions (FNAL, TJNAF, and BNL) has operated computing equipment for LQCD computing for more than 10 years.**Status:** No change

The DOE held a cost and schedule review of the project on May 24-25, 2005, and found that the cost projections for the hardware were reasonable. Operations costs are monitored and reported quarterly. In FY06, operational costs for the project stayed within budget, and the planned scientific computing was delivered on schedule. Subsequent annual DOE project progress reviews have found that cost projections for future hardware procurements and operations support were reasonable.

Risk #4: Technical Obsolescence**Date Identified:** 7/1/2004 (Modified 11/1/05 and 5/25/06)**Area of Risk:** Technology**Description:** The hardware acquired by this investment becomes obsolete before the end of the planned operations and so does not deliver scientific computing for LQCD calculations in a cost-effective manner.**Probability of Occurrence:** Low**Impact of Occurrence:** Low**Risk Rating:** 0.025**Mitigation Strategy:** The approach taken is that clusters purchased by this investment will be operated for three and a half years, and subsequently retired. These assumed lifetimes are consistent with historical life cycles observed on similar hardware over the last decade. Some cluster components, such as high performance network interconnects, have longer lifetimes, and will be reused. This strategy was used successfully during the development of the SciDAC prototype clusters that are operated as part of this investment.**Status:** No change

In December 2005, a 128-node SciDAC cluster that reached three years of operation was retired at Fermilab to free facility space. In March 2006, a 128-node SciDAC cluster that reached three and a half years of operation was retired at TJNAF when it was observed to no longer be cost-effective. In August 2005, the project proposed acquiring a cluster at TJNAF that would use the same Intel processors, motherboards, and Infiniband equipment as used on the FNAL FY05 SciDAC cluster procurement. Dual-core Intel processors released subsequent to this decision appear to be superior in price/performance, and thus the planned acquisition may have reached technical obsolescence before the planned end-of-life date (3.5 years). The use of new dual-core processors, chipsets, and motherboards that had not been evaluated in prototypes presented a

different set of technical risks (reliability, failure to meet performance goals). Via the change control mechanism, the project agreed to modify procurement plans to use the dual-core processors if the SciDAC LQCD project first purchased and evaluated a substantial (70 node) prototype. In January 2006, the SciDAC LQCD project took delivery at TJNAF of a 70-node cluster built with the new Intel dual-core processors. Extensive benchmarking and system testing over the next 30 days confirmed that the new processors and motherboards were suitable. The LQCD project executed the change control and procured a 140-node cluster based on the new dual-core processors. This cluster became operational at the beginning of April 2006.

During the DOE External Project Progress Review held May 25, 2006, reviewers from LLNL and LLBL felt that emerging Opteron “Socket-F” processors would offer superior price/performance to the Opteron dual-core processors selected by the project. The “Socket-F” processors were not available on the market but were projected to be available in time for the FY06 procurement to experience minimal delays. In response to the risk raised by the review panel that the current Opteron “Socket 940” processors could be rendered obsolete by the new “Socket-F” processors, the project asked vendors to submit additional bids based on Opteron “Socket-F” processors. Also, the project measured the performance of prototype Socket-F systems. In June 2006, new bids were received from vendors. Combined with the measured performance on Socket-F systems, the price/performance of the proposed computers was inferior to the currently available dual-core processors. Further, there was substantial risk of a delay in delivery of Socket-F systems, since the processors were not scheduled for general availability until at least September 2006. The project awarded the FY06 FNAL purchase to a vendor supplying the current dual-core Opteron processors. Delivery of the Opteron-based cluster was completed by September 2006, with full deployment achieved by December 2006.

Risk #5: Failure of Technology to Advance at Anticipated Rates

Date Identified: 7/1/2004, modified 12/20/05

Area of Risk: Technology

Description: The performance of commodity hardware components may not improve or their price may not drop as rapidly as anticipated, resulting in the investment failing to meet performance goals in the later years of the project.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.025

Mitigation Strategy: In any year, this risk is low for the current budget year; however, the risk increases when planning for succeeding years. The strategy is to follow the market carefully and build prototypes before developing large production machines. Components of clusters are carefully selected for cost effectiveness. Thus, if the network performance does not improve as expected, money can be saved on nodes by selecting slower, more cost effective CPUs whose speed will not be wasted because the network limits overall performance. This savings on each node will enable purchasing a larger number of nodes. Performance goals are set more conservatively for the later years in the project to account for market evolution uncertainty.

Status: No change

In May 2005, Fermilab brought online an Infiniband-based cluster whose price /performance was better than planned. In April 2006, TJNAF brought online an Infiniband-based cluster whose price/performance was again better than planned (original baseline: 0.2 Tflops/s, performance as

delivered: 0.32 Tflops/s). In September 2006, FNAL will bring online an Infiniband-based cluster with performance again exceeding the planned level (original baseline: 1.8 Tflops/s, expected performance: 2.3 Tflops/s). The LQCD project, in plans approved in August 2005, had planned to use new Intel processors based upon the Fully Buffered DIMM (FB-DIMM) architecture in the FY06 major procurement at Fermilab. Testing of an Intel prototype in December 2005 showed much lower performance on LQCD codes than would be required to meet FY06 milestones. To mitigate this risk of a performance shortfall, the FY06 FNAL procurement specifications were widened to allow vendors to bid dual core Opteron as well as single processor Intel systems (“Extreme Edition”). Based on evaluations of prototypes, in May a recommendation for procurement award based on AMD Opteron processors was made by the project. Intel FB-DIMM and single processor designs did not meet technical price/performance goals; designs based on AMD Opteron processors met these goals. Delivery of the Opteron-based cluster was completed by September 2006.

Risk #6: Marginal System Reliability due to System Complexity

Date Identified: 7/1/2004

Area of Risk: Reliability of Systems

Description: Complex multi-processor systems fail more frequently as they grow in size, leading to failure of the project to meet technical performance goals (e.g., delivery of computing capability).

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: Clusters of size 500-1200 processors are planned in each year of this investment. Experience gained during the SciDAC Lattice Gauge Computing Project with machines with 128 to 256 processors indicated that proposed machines of this size will run reliably.

Status: No change

In December 2005, Fermilab completed the purchase and integration of a 520-node Infiniband cluster. Operational experience with this cluster confirmed that the 1000- processor cluster planned for FY06 would be reliable.

Risk #7: Reduction in Resources Provided by Host Laboratories

Date Identified: 7/1/2004

Area of Risk: Organizational and Change Management

Description: The three host institutions (FNAL, TJNAF, and BNL) will provide space, power, network connectivity, and mass storage for the LQCD systems purchased and operated by this investment. Failure to provide any of these resources would lead to the project not meeting technical performance goals.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.025

Mitigation Strategy: The project will maintain Memoranda of Understanding (MOU) with each host institution that defines the resources that each institution will provide, and the required computer room space will be made available as agreed-to. Only a small fraction of the Internet bandwidth and mass storage of the laboratories will be required to support the LQCD project.

The experiments that are the main users of these computer facilities are a high priority for each of the laboratories, and the computer space, and network and mass storage resources will continue to evolve to support these experiments in a way that will also meet the needs of this investment. However, during the course of the project, should one of the three host institutions indicate that it will not be able to provide the required resources in a later year; the project will plan to shift deployment of hardware to one of the other host institutions.

Status: No change

At the May 2005 and May 2006 project reviews, the space, power, storage, and network requirements were presented. The reviewers approved of these plans. Additionally, Memoranda of Understanding are in place with each host institution that cover the remaining years of the investment.

Risk #8: Disruption of Operations due to Natural Disasters or Major Electrical Failures

Date Identified: 7/1/2004

Area of Risk: Surety (Asset Protection) Considerations

Description: Natural disaster and/or major electrical failure lead to disruption of operations and the failure of the project to meet technical performance goals.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: The deployment of SciDAC LQCD libraries at each site allows end users to shift their scientific production easily from one host institution to another. Should a major disruption occur, critical scientific production (as determined by the Scientific Program Committee and the LQCD Executive Committee) could continue by such a shift. This would require other less important production to be slowed or delayed. Note that no mitigation strategy is available which could sustain the rate of computations should one of the facilities suffer a major outage.

Status: No change

No major disruptions have occurred as of the date of this document.

Risk #9: Project Impact on Commodity Markets

Date Identified: 7/1/2004

Area of Risk: Strategic

Description: The Lattice QCD community becomes such a large purchaser of components that it affects the commodity market.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: Given the small size of this effort (\$9.2M over 4 years) compared to the commodity market (hundreds of billions of dollars), this appears to be an insignificant risk. No mitigation is necessary.

Status: No change

Risk #10: Changes in Agency Personnel**Date Identified:** 7/1/2004**Area of Risk:** Capability of Agency to Manage the Investment**Description:** Changes occur in DOE personnel, which affect continuity and support for this investment.**Probability of Occurrence:** Low**Impact of Occurrence:** Moderate**Risk Rating:** 0.125**Mitigation Strategy:** DOE staff has knowledge of the investment and have been providing support for over five years. As the investment spans multiple programs, this expertise is not limited to a single individual, and so the impact of a single change is minimal. The existence of an Integrated Project Team, whose composition includes Federal personnel, also mitigates risks due to agency personnel changes. A rigorous review process will be established and maintained to mitigate risks, including monthly and quarterly reports and annual reviews.**Status: No change**

Quarterly project reports began at the start of the investment (October 1, 2005). Annual reviews are completed by June 30 of each year of the investment. The first annual project progress review was held May 25, 2006. The Federal Project Manager changed in June 2006 with no disruption to the investment.

Risk #11: Failure of New Systems to Meet Goals**Date Identified:** 7/1/2004**Area of Risk:** Overall Risk of Investment Failure**Description:** A major system, such as a new cluster or a high performance network, simply fails to work and the investment does not meet technical goals.**Probability of Occurrence:** Low**Impact of Occurrence:** Severe**Risk Rating:** 0.125**Mitigation Strategy:** On an annual basis, the project evaluates prototype machines before procuring and installing production hardware. The project also builds appropriate acceptance criteria into major purchases.**Status: No change**

Plans for each year of the investment are altered according to the results of prototyping and operational experience. Annual reviews will be completed by June 30 of each year of the investment, and will verify that appropriate evaluations are being conducted. In FY06, evaluations of prototypes lead to technical design changes (at TJNAF, a dual-core processor was selected instead of the planned single core processor; at FNAL, a change to AMD processors from Intel processors occurred because of the failure of Intel systems to meet price/performance goals).

Risk #12: Adverse Impact due to Changes in Technology or Project Staff**Date Identified:** 7/1/2004**Area of Risk:** Organizational and Change Management**Description:** Changes in technology and staff can have adverse effects on the project.**Probability of Occurrence:** Medium**Impact of Occurrence:** Moderate

Risk Rating: 0.25

Mitigation Strategy: Project personnel continually study and understand changes in technology that impact the investment. The project will maintain broad expertise within its staff.

Status: No change

In August 2005, plans for systems to be built in the first year of the investment were revised to ensure that Infiniband expertise at FNAL was passed on to TJNAF staff. The FY06 cluster that became operational at TJNAF in April 2006 was the successful result of this expansion of expertise with the investment staff. Further, the project performs integrated procurements when possible across at least two of the host institutions, ensuring distribution of expertise amongst the host sites.

Risk #13: Reduction in Overall Funding Allocation

Date Identified: 7/1/2004

Area of Risk: Business

Description: Changes in funding, due to alteration in administration policy, or legislative directives.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: The investment will allocate resources and build new computing capabilities on a yearly basis, so it will be possible to adjust to changing funding levels. This is particularly so because the systems are modular, so reductions in funding can be adjusted for by reducing the size of the systems. Such reductions will delay reaching computational and scientific milestones. A strategy is not available which will mitigate the loss of technical computing capability due to substantial decreases in funding.

Status: No change

The project will adjust procurements and allocations annually according to available resources. These adjustments will be one of the subjects of the annual project reviews that will be completed by June 30 of each year. In FY06, this annual review occurred May 25, 2006.

Risk #14: Loss of Stored Data

Date Identified: 7/1/2004

Area of Risk: Data/Info

Description: Loss of stored data.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: All important data sets and systems data are backed up to tape. Essential outputs (such as the computationally-expensive vacuum gauge configurations) are stored at multiple sites. Data recovery procedures are tested at each site at least annually.

Status: No change

In FY05, the project established procedures for users to move files between the three sites and implemented dedicated mass storage areas at FNAL and TJNAF. These procedures were implemented and tested in FY06, and remain in place and functional. Data storage and replication are monitored by the IPT.

Risk #15: Technology Fails to Meet Expectations

Date Identified: 7/1/2004

Area of Risk: Technology

Description: Commercial technology does not fulfill expectations, and in the later years of the investment the project cannot meet technical objectives.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: Test individual components, build prototypes, and perform acceptance tests.

Status: No change

Prototype clusters were built at FNAL and TJNAF in 2005; the results of this prototype work drove the formation of plans for procurements in the first year of the investment (FY06). Further prototyping in FY06 refined the choices of hardware procured at TJNAF and FNAL. On an annual basis, external panels review future plans, including projected performance and cost of the available technology.

Risk #16: Changes in SC Mission or Plans

Date Identified: 7/1/2004

Area of Risk: Strategic

Description: Changes in the mission and/or plans of the Office of Science.

Probability of Occurrence: Low

Impact of Occurrence: Low

Risk Rating: 0.025

Mitigation Strategy: The computing systems acquired by this investment for LQCD computing have a broad range of applicability in other areas of computational science and could be put into other uses.

Status: No change

Completed.

Risk #17: Inappropriate Use by Unauthorized Personnel

Date Identified: 7/1/2004

Area of Risk: Security

Description: Inappropriate use of computer resources by unauthorized personnel.

Probability of Occurrence: Low

Impact of Occurrence: Low

Risk Rating: 0.025

Mitigation Strategy: The computing hardware acquired and operated by this investment is included in enclaves at each of the three sites (FNAL, TJNAF, and BNL). These enclaves have approved C&As according to Federal guidelines (NIST, DOE). Strong authentication is required for access to the systems. The computer resources are on private networks behind these secure systems. The project will coordinate security with the host laboratories. Usage is carefully monitored and controlled by batch systems. Performance is also carefully monitored, so any unauthorized usage would be quickly noticed and terminated. On clusters, batch systems

automatically terminate user processes at the end of each job and before each new job starts up. Thus, any unauthorized process would be terminated.

Status: No change

System specific security plans for each of the three sites are in place and are maintained. Each site has been C&A'd and has a current Authority to Operate (ATO).

Risk #18: Unauthorized Access to Private Information

Date Identified: 7/1/2004

Area of Risk: Privacy

Description: Unauthorized access to computing hardware can disclose private information.

Probability of Occurrence: Low

Impact of Occurrence: Low

Risk Rating: 0.025

Mitigation Strategy: No classified information, sensitive data, or personally identifiable information is stored on the systems.

Status: No change

No privacy risks are present because the Lattice QCD systems acquired and operated by the investment contain no personally identifiable information.

Risk #19: Slow Transfer Rates Inhibit Productivity

Date Identified: 6/1/2005

Area of Risk: Data/Info

Description: Slow Internet data transfer rates between the three labs inhibits productivity

Probability of Occurrence: Low

Impact of Occurrence: Low

Risk Rating: 0.025

Mitigation Strategy: FNAL, BNL, and TJNAF network staff will tune parameters to optimize data transfers. Scientific allocations of time on the LQCD clusters will take into account the quantity of data which must be transferred between sites; if network performance would limit productivity, allocations will be made such that analysis jobs would run at the same site as data are stored (i.e., to minimize transfers).

Status: No change

In FY06, the Scientific Program Committee and the LQCD project staff distributed scientific projects across the three host institutions in a manner which minimized the requirements on network data transfers between the sites. This approach will be continued for the remainder of the project.

Risk #20: Differing Authentication Schemes Limit Productivity

Date Identified: 6/1/2005

Area of Risk: Data/Info

Description: Differing authentication schemes among the three labs makes data transfers difficult and limits productivity

Probability of Occurrence: Low

Impact of Occurrence: Low

Risk Rating: 0.025

Mitigation Strategy: BNL and TJNAF use *ssh* firewalls to secure LQCD systems, whereas FNAL uses Kerberos authentication. The SciDAC Lattice Gauge Computing project was asked to design, implement, and maintain scripts and other tools to assist users.

Status: No change

In place and on going (security policies have and continue to evolve which will inevitably result in software and configuration changes at the three laboratories). Status as of Aug 2005: tools are in place and tested which allow facile movement of data by users while continuing to fulfill the computer security requirements (strong authentication) at each lab.

The SciDAC Lattice Gauge Computing project's Software Committee will monitor and report on this issue at each collaboration meeting; the last such meeting occurred in April 2006. Also in 2006, tools at TJNAF were upgraded to allow facile access and data transfers to/from FNAL systems.

Risk #21: Substantial Increase in Costs due to Increased Power Consumption

Date Identified: 8/8/2005

Area of Risk: Costs.

Description: In the experience of DOE personnel at LLNL and LBL, power consumption of computers has increased exponentially, following a Moore's Law behavior. The direct (electricity for computers) and indirect (electricity for cooling the computers) costs to the DOE could be substantial in the later years of the project.

Probability of Occurrence: Medium

Impact of Occurrence: Moderate

Risk Rating: 0.25

Mitigation Strategy: The project will work to understand historical power trends and use these trends to predict electrical costs. The project will track actual power consumption of new systems.

Status: In August 2005, the project reported on the power consumption of clusters used for LQCD calculations over the previous six years. The power consumption trend was linear, not exponential, and the costs to the host institutions during the life of the project were predicted according to this linear trend. In May 2006, the investment awarded the major system procurement for FNAL to a vendor who proposed AMD multi-core Opteron based hardware. The power consumption for this new hardware fell considerably below the linear trend observed in August 2005. The adoption of multi-core processor technology by Intel and AMD has substantially mitigated this risk of sharply increasing lifecycle costs.

Risk #22: Delayed Availability of Quad-Core Processors for the 7n Cluster

Date Identified: 7/7/2007

Area of Risk: Schedule

Description: The JLab 7n cluster was designed to use AMD quad-core processors. These processors are necessary in order to meet the performance target (sustained TFlops) for this cluster. The release schedule of the quad core processors is subject to change, and a late release by AMD will delay full implementation of the 7n cluster.

Probability of Occurrence: Medium

Impact of Occurrence: Moderate

Risk Rating: 0.25

Mitigation Strategy: The contract used to procure 7n specified that the hardware was to be delivered with currently available dual-core AMD processors, with an option to upgrade to quad-core processors when they became available. With dual-core processors, the 7n cluster delivers approximately 60% of the performance target. In June 2007, the project declared a deadline of September 15 for the decision of exercising the quad-core option. By this date, if the project has successfully tested quad-core chips on a portion of the 7n cluster, where success means that the quad-core chips will deliver the target performance without compromising operability or usability, the option to replace the dual-core chips with quad-cores will be exercised. Otherwise the project will purchase additional dual-core systems with the funds that had been held in reserve for the processor upgrade. Note that if the option is exercised, there may still be a delay in upgrading the cluster because of limited availability of the new processors. During this delay, the cluster will still be able to deliver TFlops for LQCD calculations.

Status: New

Sample quad core processors from AMD are expected for testing by the middle of August 2007. These processors will be a slower speed than those specified in the option clause of the contract; however, they will be sufficient to demonstrate operability and usability. Based on benchmarks on LQCD code obtained in May 2007 using such lower speed processors, the upgrade option would enable 7n to meet the performance target even at the lower clock speed (1.8 GHz vs 2.1 GHz).

Risk #23:

Date Identified: 7/7/2007

Area of Risk: Schedule

Description: The final procurement of the project will be a cluster purchase by Fermilab, deployed at the end of calendar year 2008, using FY08 and FY09 project DME funds. In order to meet the performance target for this cluster, new processors and chipsets from either Intel or AMD will be used. However, these processors and chipsets will not be available for evaluation until mid-FY08. There are both schedule and performance risks: the new chips may not provide sufficient performance to meet the target and the new chips may not be available in time for deployment in late calendar 2008.

Probability of Occurrence: Medium

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: Project personnel will continue to work with Intel and AMD to understand the new technologies that will be used in the new processors and chipsets. Further, testing of pre-release hardware will occur as soon as the vendors can make hardware available. If vendor roadmaps and/or pre-release hardware evaluations indicate schedule or performance slip, the project plan will be modified according to the documented change control process.

Status: New

As of July 13, 2007, vendor roadmaps indicate that the first of two new Intel processor and chipset combinations will be available to the project for evaluation in mid-August 2007, and that the second set will be available in early calendar 2008.

Risk #24: Failure to Secure Follow-on Funds to Operate Cluster through End-of-Life

Date Identified: 7/7/2007

Area of Risk: Strategic

Description: The DOE (or OMB) could decide that it will not fund a follow-on project to the current project (SC LQCD). At the end of SC LQCD in Sept 2009, the clusters purchased by the project and deployed at FNAL and JLab will be in mid-life. There is a risk that these clusters would have to be shut down or otherwise made unavailable to the US LQCD community unless funds for their steady state operation can be made available.

Probability of Occurrence: Low

Impact of Occurrence: Moderate

Risk Rating: 0.125

Mitigation Strategy: In the event that a follow-on project proposal is rejected, the US LQCD community will request steady state operations funding via a de-scoped project proposal. If this does not succeed, the community will negotiate with the host laboratories (FNAL and JLab) for operations funding from their base budgets.

Status: New

During the spring and summer of 2007, the US LQCD community has been in active communications with DOE SC about the follow-on project.

Appendix A provides a summary of the risk ratings for the LQCD project. As predicted in the initial project planning, cost and schedule elements have the highest risk ratings, with technology having the third highest rating.

7. RISK MONITORING

The risk register is reviewed and updated regularly. During a given year, whenever the probability of occurrence and impact of occurrence of an individual risk changes, the status change is indicated in the register. The status of the new risks added to the register in a given year is identified as new. New and revised mitigation actions are also recorded. The LQCD Risk Management Plan is updated annually.

8. REFERENCES

- Lattice QCD Project Execution Plan, December 2005
- Lattice QCD Work Breakdown Structure (Baselined in August 2005)
- A Guide to the Project Management Body of Knowledge (PMBOK® Third Edition), Project Management Institute
- OMB Circular Number A-11 Part 7 Capital Programming Guide V2.0 (2006) Appendix 5.

APPENDIX A: SUMMARY OF RISK RATINGS BY CATEGORY

| ID | Business | Cost | Data | Org./ Change | Privacy | Project Mgmt. | Schedule | Security | Strategic | Technology |
|------------------------|--------------|--------------|--------------|-----------------|--------------|------------------|--------------|--------------|--------------|--------------|
| 1 | | | | | | | 0.25 | | | |
| 2 | | 0.25 | | | | | | | | |
| 3 | | 0.125 | | | | | | | | |
| 4 | | | | | | | | | | 0.025 |
| 5 | | | | | | | | | | 0.025 |
| 6 | | | | | | | | | | 0.125 |
| 7 | | | | 0.025 | | | | | | |
| 8 | | | | | | | | 0.125 | | |
| 9 | | | | | | | | | 0.125 | |
| 10 | | | | | | 0.125 | | | | |
| 11 | | | | | | | | | | 0.125 |
| 12 | | | | 0.25 | | | | | | |
| 13 | 0.125 | | | | | | | | | |
| 14 | | | 0.125 | | | | | | | |
| 15 | | | | | | | | | | 0.125 |
| 16 | | | | | | | | | 0.025 | |
| 17 | | | | | | | | 0.025 | | |
| 18 | | | | | 0.025 | | | | | |
| 19 | | | 0.025 | | | | | | | |
| 20 | | | 0.025 | | | | | | | |
| 21 | | 0.25 | | | | | | | | |
| 22 | | | | | | | 0.25 | | | |
| 23 | | | | | | | 0.125 | | | |
| 24 | | | | | | | | | 0.125 | |
| Grand Total | 0.125 | 0.625 | 0.175 | 0.275 | 0.025 | 0.125 | 0.625 | 0.150 | 0.275 | 0.425 |

APPENDIX B: GLOSSARY

| Acronym | Definition |
|------------|---|
| 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 | Center Processor Unit |
| DOE | Department of Energy |
| FNAL | Fermi National Accelerator Laboratory |
| Intel | A processor company |
| IPT | Integrated Project Team |
| LBL | Lawrence Berkeley National Laboratory |
| LLNL | Lawrence Livermore National Laboratory |
| LQCD | Lattice Quantum Chromodynamics |
| NCSA | National Center for Supercomputing Applications |
| NERSC | National Energy Research Scientific Computing Center |
| QCD | Quantum Chromodynamics |
| QCDOC | QCD On a Chip (BNL Supercomputer) |
| SciDAC | Scientific Discovery through Advanced Computing |
| TFlop/s | Teraflops per second, 1 teraflop = 10^{12} flops |
| TFlop/s-yr | Computing delivered by 1 TFlop/s sustained for one year |
| TJNAF | Thomas Jefferson National Accelerator Facility |