LQCD-ext Risk Register

Version	Date	Description of Change
1	8/18/2009	Initial Risk Items for LQCD-ext (derived from LQCD project)
2	3/16/2010	Revised Risk Mitigation Strategies
3	7/21/2010	Revised Risk Management Plan V1.2
4	4/26/2011	Revised Risk Register for GPU/Ds extension purchase

LQCDEXT PROJECT RISK REGISTER

ID FY11	Probability of Occurrence (Initial)	FY11 Risk Area	Description	Probability of Occurrence (FY11)	Impact of Occurrence (FY11)	Risk Rating	Last Status	Last Change	Intial Date	Last Update
1	Low	Schedule	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; c) Funding may be lower than anticipated.	Medium	Moderate	0.25	Open	No change	7/1/04	7/21/09
2	Low	Cost	Although cost projections for the current budget year are reasonably precise, projections for subsequent years become progressively uncertain.	Medium	Moderate	0.25	Open	No change	7/1/04	7/21/09
3	Low	Cost	Unexpected increases in life costs arise after systems are acquired.	Low	Moderate	0.125	Open	No change	7/1/04	7/21/09
4		Technology	Obsolecence: 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.	Low	Low	0.025	Open	No change	7/1/04	7/21/09
5		Technology	Feasibility: 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.	Low	Low	0.025	Open	No change	7/1/04	7/21/09
7	Low	Cost	Dependency: Host institutions will not provide space, network connectivity, and mass storage.	Low	Moderate	0.125	Open	Changed ratings	7/1/04	7/21/09
8	Low	Security Safeguard	Surity: A major failure of a facility due to natural disaster (destruction of buildings, utility systems)	Low	Severe	0.125	Open	Split into two risks	7/1/04	7/21/09
10		Cost	Agency personnel changes, limiting continuity and support for this investment.	Low	Moderate	0.125	Open	No change	7/1/04	7/21/09
11	Low	Technology	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.	Low	Severe	0.125	Open	Changed description	7/1/04	7/21/09
12		Technology	Performance: Changes in technology and staff can have adverse effects on the project.	Medium	Moderate	0.25	Open	No change	7/1/04	3/16/10

LQCDEXT PROJECT RISK REGISTER

ID FY11	Probability of Occurrence (Initial)	FY11 Risk Area	Description	Probability of Occurrence (FY11)	Impact of Occurrence (FY11)	Risk Rating	Last Status	Last Change	Intial Date	Last Update
13	Medium	Cost	Changes in funding, due to alteration in administration policy, or legislative directives.	Low	Moderate	0.125	Open	No change	7/1/04	7/21/09
16		Technology	Changes in the mission and plans of the Office of Science.	Low	Low	0.025	Open	No change	7/1/04	7/21/09
17		Security Safeguard	Inappropriate use of computer resources by authorized or unauthorized personnel	Low	Moderate	0.125	Open	Changed description	7/104	7/21/09
18		Security Safeguard	Unauthorized access to computing hardware can disclose private information.	Low	Low	0.025	Open	No change	6/1/05	7/21/09
19		Technology	Slow Internet data transfer rates among the three labs and external sites may inhibit productivity	Low	Low	0.025	Open	No change	6/1/05	7/21/09
20		Technology	Differing authentication schemes among the three labs makes data transfers difficult which limits productivity	Low	Low	0.025	Open	No change	6/1/05	3/16/10
21		Cost	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.	Medium	Moderate	0.25	Open	No change	8/8/05	9/25/06
24		Cost	Risk of unavailability of DOE funding beyond the end of the project (end of FY09)	Low	Moderate	0.125	Open	No change	7/7/07	7/7/07
25		Technology	Complex multi-processor systems may not perform adequately due to unforseen bottlenecks as core counts rise that are not addressed adequately in software, leading to failure of the project to meet technical performance goals (delivery of computing capability and/or capacity	Low	Moderate	0.125	Open	New	7/21/09	7/21/09
26	Low	Security Safeguard	Surity: Utility system failure at one of the facilities	Medium	Moderate	0.25	Open	New: Split from another risk	7/21/09	7/21/09
27	Low	Security Safeguard	Reliability: Loss of nearline stored data.	Low	Moderate	0.125	Open	New: Split from another risk item	7/1/04	3/16/10

LQCDEXT PROJECT RISK REGISTER

ID FY11	Probability of Occurrence (Initial)	FY11 Risk Area	Description	Probability of Occurrence (FY11)	Impact of Occurrence (FY11)	Risk Rating	Last Status	Last Change	Intial Date	Last Update
28		Security Safeguard	Data Integrity: Some stored data may get corrupted or lost. Some LQCD data products, such as gauge configurations and very large quark propagators, are very valuable in terms of the computing required to reproduce them in case of loss or corruption.	Low	Low	0.025	Open	New	8/18/09	3/16/10

LQCDEXT RISK MITIGATION STRATEGIES

ID FY11	Description	Mitigation Strategy	Risk Response Action	Date of Response Action Taken
1	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; c) Funding may be lower than anticipated.	Over the past five years, the LQCDEXT investment team worked on multiple large cluster hardware procurements using DOE LQCD project (FY06-FY09) and the DOE SciDAC Lattice QCD Computing Project funds with significant success. Experienced professional staff follows the commodity market carefully and gains insight by evaluating prototype hardware. They meet with vendors frequently under non-disclosure agreement 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, the team has the capability of testing prerelease components. Working with the manufacturers the team is aware of deficiencies in vendor products. The team is able to determine whether new capabilities will actually provide any advantage in future system procurements. The team plans to use past procurement methodologies fine tuning them as appropriate.		
2	Although cost projections for the current budget year are reasonably precise, projections for subsequent years become progressively uncertain.	Market information is gathered and prototypes are built throughout the lifetime of the project. Open procurements of commodity components allow for competitive prices. Since hardware is modular in nature, if prices exceed expectations in any given year, it is possible to deploy smaller machines. . A level of performance contingencies are maintained for all procurements.		
3	Unexpected increases in life costs arise after systems are acquired.	Hardware maintenance costs are included in procurement of components for each new system procured (each year). Operations costs are well understood based on years of similar operational experience. Each of the three host institutions (FNAL, TJNAF, and BNL) has operated computing equipment for LQCD computing for more than 10 years. Since the LQCD project is staffed by few key professionals, the loss of any of them is likely to affect the performance of the project; this risk is accepted “as-is” although the project does strive through cross-training and other efforts to maintain expertise across and among the staffs at the three sites.		
4	Obsolecence: 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.	Clusters purchased by this investment are 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.		
5	Feasibility: 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.	In any year this risk is low for the current budget year since the price/performance ratio is well defined for the current year. However, the risk increases when planning for the succeeding year. 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.		
6	Complex multi-processor systems fail more frequently as they grow in size, leading to failure of the project to meet technical performance goals (delivery of computing capability).	Closed		
7	Dependency: Host institutions will not provide space, network connectivity, and mass storage.	The required computer room space is available at each of the host institutions. Only a small fraction of the Internet bandwidth and mass storage of the laboratories is required to support the LQCDEXT project. The experiments that are the main users of 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. Further, the project maintains Memoranda of Understanding (MOU) with each institution which detail the resources which are to be committed. In any given year, should one of the three host institutions predict that it would 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.		

ID FY11	Description	Mitigation Strategy	Risk Response Action	Date of Response Action Taken
8	Surity: A major failure of a facility due to natural disaster (destruction of buildings, utility systems)	LQCD computer facilities are located within large buildings suitable for large computing installations. These building are not necessarily hardened for natural disasters. To make them disaster-proof would be extremely expensive. The impact of a disaster is severe because this will impact the scientific delivery schedule significantly. However, the probability of occurrence is low. The project accepts this risk.		
9	Monopoly: Community becomes such a large purchaser of components that it effects the market for them.	Closed	Closed	7/21/2009
10	Agency personnel changes, limiting continuity and support for this investment.	DOE staff has knowledge of the investment, and have been providing support for over six 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 mitigate risks due to agency personnel changes. A rigorous review process has been established to mitigate risks, including monthly and quarterly reports and annual reviews.		
11	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.	The project evaluates prototype machines before procuring and installing production hardware (annually). The project also builds appropriate acceptance criteria into major purchases. During the acceptance testing phase lasting 30 days, the system is tested thoroughly. If the system is deemed to be unacceptable, it can be returned to the supplier under the warranty condition. The project procures systems with a minimum 3 year warranty service, 4 hour response, 48 hour repair service response.		
12	Performance: Changes in technology and staff can have adverse effects on the project.	Project personnel continually study and understand changes in technology that impact the investment. The project maintains a broad range of expertise within its staff.		
13	Changes in funding, due to alteration in administration policy, or legislative directives.	The investment allocates resources and builds new computing capabilities on a yearly basis, so it is 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 may delay reaching computational and scientific milestones. A strategy is not available which mitigates the loss of technical computing capability due to substantial decreases in funding.		
14	Loss of archival stored data.	Closed		
15	Commercial technology does not fulfill expectations, and in the later years of the investment the project cannot meet technical objectives	Based on the past experience of the project, commercial technology has fulfilled the expectations of the project. During the history of the project, this was never a problem. However, the project personnel continue to pursue comprehensive benchmarking and testing of individual components, building prototypes, and performing acceptance tests.		
16	Changes in the mission and plans of the Office of Science.	The computing systems acquired by this investment for LQCDEXT computing have a broad range of applicability in other areas of computational science and could be put into other scientific uses. This is an accepted "as-is" risk.		
17	Inappropriate use of computer resources by authorized or unauthorized personnel	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.		
18	Unauthorized access to computing hardware can disclose private information.	No classified information, sensitive data, or personally identifiable information is stored on the systems. No privacy risks are present because the lattice QCD systems acquired and operated by the investment contain no personally identifiable information. To enforce this, LQCD users are required to comply with security policies established by respective laboratories.		

ID FY11	Description	Mitigation Strategy	Risk Response Action	Date of Response Action Taken
19	Slow Internet data transfer rates among the three labs and external sites may inhibit productivity	FNAL, BNL, and TJNAF network staff tunes parameters to optimize transfers. Scientific allocations of time on the LQCDEXT clusters takes into account the quantity of data which must be transferred between sites; if network performance would limit productivity, allocations are made such that analysis jobs would run at the same site as data are stored (i.e., to minimize transfers). This is an accepted risk for the project since (~~controls for computer security protections are expected to become stricter in near future.~~) the data transfer rates or available bandwidth will not keep up with the amount of data to be transferred.		
20	Differing authentication schemes among the three labs makes data transfers difficult which limits productivity	FNAL, BNL, and TJNAF network staff tunes parameters to optimize transfers. Scientific allocations of time on the LQCDEXT clusters takes into account the quantity of data which must be transferred between sites; if network performance would limit productivity, allocations are made such that analysis jobs would run at the same site as data are stored (i.e., to minimize transfers). This is an accepted risk for the project since controls for computer security protections are expected to become stricter in near future.		
21	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.	project staff uses historical power trends to predict electrical costs. The project also tracks actual power consumption of new systems. The project also specifies power consumption criteria for new procurements to prefer lower power components. The project is always investigating new cost saving and effective computer cooling technologies.		
22	Delay in the release of AMD Quad-processors for Jlan 7n cluster deployment	Closed		
23	Schedule concern for the processor & chipset delivery for FNAL FY08 cluster deployment	Closed		
24	Risk of unavailability of DOE funding beyond the end of the project (end of FY09)	Closed		
25	Complex multi-processor systems may not perform adequately due to unforseen bottlenecks as core counts rise that are not addressed adequately in software, leading to failure of the project to meet technical performance goals (delivery of computing capability and/or capacity	LQCD project has been using multi-processor systems for a while now without experiencing any major software issues. However, there is a likelihood that the LQCD software may come across some issues with multiprocessor systems. The LQCD staff and the off-project LQCD software development team is watching for any such possibilities taking various actions as necessary.		
26	Surity: Utility system failure at one of the facilities	There is a moderate possibility of a single-site utility failure. However, the deployment of SciDAC LQCDEXT libraries at each site allows end users to shift their scientific production easily from one host institution to another. Should a significant disruption occur, critical scientific production (as determined by the Scientific Program Committee and the Lattice QCD Executive Committee) could continue by such a shift. This may require other less important production to be slowed or delayed. Note that no mitigation strategy is available which could sustain the normal rate of computations should one of the facilities suffer a major utility outage.		
27	Reliability: Loss of nearline stored data.	The LQCD project makes every effort to provide adequate near-line storage to run the simulation jobs. This includes Lustre based storage at FNAL and NFS based storage at TJNAF. Related procedures and technologies are refined continuously. Currently, the project has more than ade-quate near-line storage. A formal decision has been made that LQCD project is not responsible for the archival storage data.		
28	Data Integrity: Some stored data may get corrupted or lost. Some LQCD data products, such as gauge configurations and very large quark propagators, are very valuable in terms of the computing required to reproduce them in case of loss or corruption.	The most precious LQCD data products (i.e., the most expensive to reproduce) are gauge configurations. By USQCD policy, overseen by the Executive Committee, to prevent against loss these configurations are stored on tape at two or more geographically diverse sites. The responsibilty for this storage is held by the individual physics collaborations that have generated the particular data ensembles. To guard against silent corruption, by policy these files must be written with checksum (32-bit CRC) data that can be compared on subsequent access to determine whether any data changes have occurred. The USQCD standard I/O library, QIO, can be used to calculate, store, and compare these CRC data. The USQCD user community are also urged in documentation and at the annual collaboration meeting to use this data integrity facility of QIO to guard quark propagator and other data products. Also, single gauge configurations can be regenerated from prior gauge configurations.		

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

Risk Matrix

0.675	0.375	0.075
0.45	0.25	0.05
0.125	0.125	0.025

Risk rating

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