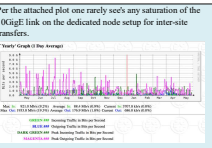


Risk ID	Risk Title	Risk Area	Description	Probability of Occurrence	Impact of Occurrence	Risk Rating	Risk Priority	Risk Status	Creation Date	Next Review	Last Change	Mitigation Strategy	7/24/18 notes and 5/28/19 notes	Notes by Aaming (FNL) 5/30/2019	Notes	
1	01: Technology/systems may take longer than expected to become available	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.	High	Moderate	0.125	3 - Low	Exists	7/1/04	5/29/19	No change	For more than a decade now, the LQCD Integrated Project Team has worked on multiple large cluster hardware procurements with significant success. Experienced professional staff follow the commodity market carefully and gain insight by evaluating prototype hardware. They meet with vendors frequently under non-disclosure agreements and are briefed on roadmaps for components such as processors, chipsets, motherboards, network interface cards and switches. In addition to working closely with manufacturers and system integrators, the team has the capability of testing pre-release components. Working with the manufacturers, the team is aware of the strengths and weaknesses 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.	FNL: Probability and impact are still high.			
2	02: Cost projections for future years uncertain	Cost	Although cost projections for the current budget year are reasonably precise, projections for subsequent years become progressively uncertain.	Medium	Moderate	0.225	2 - Medium	Exists	7/1/04	5/29/19	No change	Market information is gathered and prototypes are built throughout the lifetime of the project. Open procurements of commodity components allows 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.	FNL: With the Institutional cluster model we will continue to make projections using the same mitigation strategy.		If conventional clusters remain competitive for next two years, the risks will remain same.	
8	08: Failure of a facility due to natural disaster	Service	Sturity: A major failure of a facility due to natural disaster (destruction of buildings, utility systems)	Low	Severe	0.125	3 - Low	Exists	7/1/04	5/29/19	No change	LQCD computer facilities are located within large buildings suitable for large computing installations. These buildings 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.	FNL: No change			
10	10: Agency personnel changes reduce support for project	Cost	Agency personnel changes, limiting continuity and support for this investment.	Low	Moderate	0.050	3 - Low	Exists	7/1/04	5/29/19	No change	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.	FNL: No change			
11	11: Major computer system failure	Technology	A major system, such as a cluster or a high performance network, fails to meet performance specifications such that our ability to achieve scientific goals is compromised and the investment does not meet technical goals.	Low	Moderate	0.125	3 - Low	Exists	7/1/04	5/29/19	No change	While this risk applies to all systems in principle, we focus our mitigation strategy on new systems since, in our experience, that is where it is most likely to occur. The project evaluates prototype machines before procuring and installing production hardware. 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. Also, each project purchase represents an addition of <= 50% to the deployed CPU power, which limits the impact of this risk. Even if a new system completely fails to perform despite the aforementioned safeguards, at worst only 1/3 of the post-purchase CPU power is affected. The loss of any one resource for 2-3 months would not result in a major impact on project deliverables.	FNL: With the institutional cluster model, the risk has been passed on to the institution. The project has the flexibility to buy computing time within the 2 sites. (BNL & FNL). BNL: There are 3 distinct clusters from which LQCD can purchase compute cycles. The ratio between the 3 can be adjusted for performance and usability reasons, as needed.			
13	13: Changes in funding due to policy changes or new directives	Cost	Changes in funding, due to alteration in administration policy, or legislative directives.	Low	Moderate	0.125	3 - Low	Exists	7/1/04	5/29/19	No change	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.	FNL: With the IC model we would purchase computing cycle times to match available funding. Risk that we still have: Such reductions may delay reaching computational and scientific milestones. There is currently no strategy available that would mitigate the loss of technical computing capability.			
16	16: Change in agency mission	Technology	Changes in the mission and plans of the Office of Science.	Low	Low	0.250	2 - Medium	Exists	7/1/04	5/29/19	No change	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 scientific uses. This is an accepted "waste" risk.	FNL: No change based on Institutional Cluster model			
17	17: Inappropriate use of computer resources	Security	Inappropriate use of computer resources by authorized or unauthorized personnel	Medium	Low	0.050	3 - Low	Exists	7/1/04	5/29/19	No change	The computing hardware acquired and operated by this investment is included in enclosures at each of the three sites (FNL, TNAF, and BNL). These enclosures 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.	FNL: No change			
18	18: Unauthorized access to computing may disclose private information	Security	Unauthorized access to computing hardware can disclose private information	Low	Low	0.025	3 - Low	Exists	6/1/05	5/29/19	No change	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.	FNL: No change based on Institutional Cluster model			
19	19: Slow networking between sites inhibits productivity	Technology	Slow Internet data transfer rates among the three labs and external sites may inhibit productivity	Low	Low	0.050	3 - Low	Exists	6/1/05	5/29/19	On 5/29 FNL updated mitigation strategy 4/22/2015: Revote mitigation strategy to address improvements in networking in past few years.	At FNL, a dedicated node to be used for inter-site transfers (via Globus Online) was deployed in 2013 with 10 GbE connectivity to the internet and QDR InfiniBand connectivity to the FNL LQCD Lustre filesystem. When users report slow transfers, Fermilab networking staff have worked with external sites (for example, Globus Online, ANL/ANHEP) to determine and repair any bottlenecks. Similarly, JLab has a dedicated 10 GbE / 40G IB data gateway hosting GlobusOnline, with (shared) 10 GbE to EShet; network experts work with EShet to diagnose any slow connections.	FNL: Same however would like to find out BNL status. BNL: two dedicated nodes each with 2x10GigE connectivity to the internet. BNL has dedicated networking staff that monitors and addresses performance issues with EShet. BNL has a Globus endpoint, but it is only connected to the home directory area on the GFPS storage system.		For the attached plot one rarely sees any saturation of the 10GigE link on the dedicated node setup for inter-site transfers.	Transfer needs between FNL and TNAF are minimal. Transfer needs between FNL and BNL are rare and the connectivity is excellent. Transfer needs between TNAF and BNL are minimal. Transfer needs are more frequent from Leadership class computing to LQCD computing sites. Although transfer rates between ANL and FNL are not an issue, there are sporadic issues with transfers from Oak Ridge to FNL. Oak Ridge to TNAF: any transfer problem occurs in bursts, mostly with propagators. Transfer problems are often solved by providing better tools to users (e.g. BHFTF, Globus Online) or suggesting procedural changes such as pre-staging from tape to disk.
20	20: Authentication differences affect inter-site transfers, productivity	Technology	Differing authentication schemes among the three labs makes data transfers difficult which limits productivity	Medium	Low	0.125	3 - Low	Exists	6/1/05	5/29/19	4/8/2015: Modest text change in mitigation strategy	FNL, BNL, and TNAF network staff tunes parameters to optimize transfers. Scientific allocations of time on the LQCD computing 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. Site Managers try to mitigate this risk by addressing helpdesk requests and better documentation.	FNL: No change			
21	21: Power costs could become substantial	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	Low	0.050	3 - Low	Exists	8/8/05	5/29/19	No change	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.	FNL: The ownership of the mitigation plan belongs to the institution not the project. BNL-MOU signed with LQCD is renewed annually and indirect costs such as power are reviewed and adjusted accordingly.			
26	26: Utility system failure at one of the facilities	Service	Utility system failure at one of the facilities	Low	Moderate	0.025	3 - Low	Exists	7/21/09	5/29/19	On 5/29/19 FNL updated Mitigation Strategy 4/8/2015: Adjusted mitigation strategy text.	There is a moderate possibility of a single-site utility failure. However, the deployment of SciDAC LQCD software libraries at each site allows end users to shift their scientific production easily from one institution to another. If 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 mitigation strategy is available which could sustain the normal rate of computations should one site facilities suffer a minor utility outage.	FNL: The summer high temp load issue has been addressed successfully. BNL: data center uses UPS-protected power backed up by generator and has redundant cooling units.		At FNL the LQCD clusters are housed in a facility that is backed up by UPS. The UPS allows sufficient time to conduct a graceful shutdown in case of an unplanned power failure. In the event of a planned power outage, facility has generator plans to fall back on generator power.	
27	27: Loss of nearline stored data	Service	Reliability: Loss of nearline stored data.	Low	Moderate	0.125	3 - Low	Exists	7/1/04	5/29/19	On 5/29/19 FNL updated Mitigation Strategy 4/8/2015: Adjusted mitigation strategy text.	The LQCD computing project makes every effort to provide adequate near-line storage to run the simulation jobs. This includes Lustre based storage at FNL and TNAF. Related procedures and technologies are refined continuously. Currently, the project has more than adequate near-line storage. A formal decision has been made that LQCD project is available for the archival storage data. The project reflects all aging storage hardware in FY13, and in FY15 will be migrating from an older Lustre v1.8 to a more stable v2.5 release.	FNL: Probability of loss (partial loss) has increased because of the aging of storage hardware at FNL. BNL operates enterprise-class GFPS storage with hardware RAID and server redundancy.		At FNL, deploying ~1TB of new Lustre storage in FY19 which will have software RAID and server redundancy. This storage hardware will be under warranty till FY23. The project (~14TB) area is backed up nightly and backs previous data such as output logs, meson correlators, and other small data files. User home areas are backed daily to the site-wide data backup facility.	
45	45: Recommendation to add the following Risk: FNL institution cluster model is new. We will need to determine if it is optimal for the LQCD community.	Technology	In order to make the operations of the FNL-IC as cost effective as possible at first order the configuration of the cluster is a consolidation of best practices learnt from administrating HPC clusters and HTC farms. There are a lot of configurations which are a "test" for the Fermilab/SLD operations team.					2-Medium	Exists	5/29/2019	Updated on 5/29/19 to show FNL NEW R Risks	For cluster configurations that are a "test" we have the option to fall back to a "known" configuration if the first attempt fails to meet expectations/needs. For e.g. instead of using an InfiniBand based router server to route in and out of public net traffic from the worker nodes, the original configuration has each worker node directly connected to the Fermilab core WAN. If there is a performance penalty or overhead as a result of this, the plan is to setup NAT (Network Address Translation) to route all such traffic via the "router host". The "router host" has been purchased but will remain inactive till such time as when we need it or if not then it will be repurposed accordingly.	Added Risk on 5/29/19			

28	28: Stored data may get corrupted or lost	Service	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	0.050	3 - Low	Exists	8/18/09	5/29/19	5/10/2017: Expanded node per Chip Watson.	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 responsibility 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 IO 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.	FNAL: Mitigation strategy is followed. Remain the same.	TINAF checks MD-5 checksum of files coming back from tape (and for raw data from experimental program calculates it soon after it is written to disk in the counting house). These checks insure that retrieved data is never corrupted by the tape library. JLab is also using OpenPFS risk-2% read check on every read, ensuring that any single disk error is never undetected and is fact corrected on the fly. The minor performance impact is accepted for the sake of higher data reliability.
37	37: Staff changes have adverse effect	Technology	Performance: Changes in staff can have adverse effects on the project.	Medium	Moderate	0.375	2 - Medium	Exists	7/1/04	5/29/19	4/22/2015: Set to impact to Low, add Notes about variants.	The project maintains staff depth in key roles: Project Manager, BNL Site Manager, FNAL Site Manager, and TINAF Site Manager. For these roles, an active deputy exists who can fill the role if and when necessary. This should keep the impact of any one key staff member Low, assuming we lose only one key staff member within a period of 6 months.	FNAL: High Risk as Amtoj is Site Mgr. and Site Arch. BNL: Low-risk since multiple facility staff possess overlapping expertise, and loss can be mitigated.	While the impact of losing more key staff will be greater, the probability of this happening within a 6 month timeframe is lesser. The case of losing two key staff members for example might have a Moderate impact instead of Low, but the probability of this is considered Low instead of Medium. In either case, the Risk Priority is the same.
38	38: Inaccurate Storage Forecasting	Cost	Changes in science algorithms or storage use patterns could lead to underestimation of future storage needs which drive up costs or limit the science that can be done with the deployed CPU's.	Low	Low	0.250	2 - Medium	Exists	8/20/14	5/29/19	5/29/19 all 3 sites have a documented Data Management plan. Changing to Low Probability	Annual review of storage needs and use patterns. Continue to employ storage "costing" in the allocation model to encourage efficient use of storage, as is done with CPU time, without negatively impacting science production. Discourage storage use not directly related to USQCD science goals. Enforce a coherent storage management policy on all sites.	FNAL: Have placed storage limits to mitigate risk. Suggest that a Data management plan is created. ACPM has scheduled meeting to review storage issues. BNL: storage quotas are in place and can be adjusted as needed. Possible to archive infrequently used data to tape to free on-disk space.	
39	39: Inadequate Lustrre Support	Cost	Lustrre may require more site effort than currently expected (for a fixed amount of storage) if vendor support or storage operating characteristics change.	Medium	Low	0.125	3 - Low	Exists	8/20/14	5/29/19	No change	Annual review of effort expended in Lustrre support and revision of forecasted support effort level. In FY15 in particular, we plan to upgrade Lustrre systems to v2.5, which could spike the support effort required. We will track the upgrade effort expected required to determine whether fallback plans are required. We can delay one or both site upgrades, slow one or both site upgrades to reduce effort expended to a tolerable level, or stagger the upgrades across the sites. The real impact is to draw personnel away from other tasks, thus degrading operations, which could have a scientific impact.	FNAL: Mitigation strategy is outdated. Intel has open sourced Lustrre major and minor versions and the open SFS. Open Source has more control with future Lustrre support and development. Sufficient support for Lustrre.	Mitigation Strategy remains the same.
41	41: Software infrastructure may not be mature enough for newer computing architectures	Technology	41: Software infrastructure may not be mature enough for the latest highest-performing architectures to allow the project to exploit the otherwise most cost-effective hardware	Low	Moderate	0.250	2 - Medium	Exists	4/22/15	5/29/19	No change	It is not necessary to suddenly have 100% of our software able to absorb a new architecture, as we are always running machines as much as 4 years older. Thus, in a single year, the newest machine might be only 25% - 40% of the total project capacity. In each procurement, we optimize the old/new machines to maximize science across a portfolio of applications (some more mature than others with respect to newer hardware). The computing project does not develop application software, and so can only interact with the complementary projects to attempt to optimize the science output across all relevant projects and machines.	FNAL: With the Institutional cluster model, the risk can be reduced to "low".	The new Xeon Phi coupled to the new Intel Omni Path fabric has challenges in achieving high bandwidth utilization. This was a known problem at procurement, and is addressed by using multiple threads to do the communication. This requires an evolution in software, and most recent codes are achieving good bandwidth, but there are still many codes that will need changes to run on large number of nodes.
42	42: LQCD sustained performance-per-dollar on commodity hardware may not improve as rapidly as anticipated	Technology	The performance of commodity hardware components may not improve or their prices may not drop as rapidly as anticipated. New or advancing technologies in commodity hardware, particularly in processors and accelerators, may not perform adequately because of unforeseen bottlenecks that are not adequately addressed in current LQCD software. Realization of these risks may result in the investment failing to meet performance goals in the later years of the project.	Medium	Moderate	0.125	3 - Low	Exists	8/19/15	5/29/19	No change	At any given time, this risk is low for the current budget year since price/performance is reasonably well known for the coming year. Since there are 2 current advanced architectures and possibly more in the near future, our mitigation is to look for where the software portfolio can get the best gain, while also feeding technology trends back into the software development projects so that they are aware of changes that need to be absorbed.	FNAL: Remain the same with medium risk	Conventional processors are no longer doubling in performance per dollar even on a 24 month schedule, and thus the project is dependent on an increasing percentage of our portfolio being able to move to advanced architectures in order to see the assumed 24 month doubling of the project's portfolio's performance per dollar.
44	44: DOE funding unavailable beyond FY19-FY23	Cost	Risk of unavailability of DOE funding beyond the end of the program (end of FY23)	High	Severe	0.050	3 - Low	Exists	10/24/16	5/29/19	No change	The project must accept this risk. Since we will not know the future project/extension funding decision until the current year funds have been committed, we will address this by adding some flexibility to the acquisitions in case some funds need to be held for operations of existing facilities to their end-of-life. We will always prepare and circulate a few alternative scenarios to possibly modify the current funding profile and/or acquisition plans in case DOE funding is not available.	Final: Institutional cluster model has been agreed upon, and we are buying yearly time instead of purchasing hardware.	During CR16-01.02 discussion, the CCB recommended preparing a few alternative scenarios for funding post-FY19 operations with project funds. 3/15/2017 Status: USQCD (Paul M) to present scenarios at upcoming DOE Review. 9/18/17 (BB) - Future funding is uncertain based on HEP position taken in response to FY18 PBR. Increased risk probability to high. Project is focusing on securing funding for FY18-19 through transition to a hybrid Institutional Cluster operating model. This risk should be reviewed in 3 mos. to re-assess situation.