

Managing Government Funded Scientific Consortia

DR. BAKUL BANERJEE, *MEMBER, IEEE*

Abstract—In recent years, it is becoming apparent that good science not only requires the talents of individual scientists, but also state-of-the-art laboratory facilities. These facilities, often costing millions to billions of dollars, allow scientists unprecedented opportunities to advance their knowledge and improve the quality of human life. To make optimum use of these experimental facilities, a significant amount of computational simulations is required. These mega-projects require large-scale computational facilities and complementary infrastructures of network and software.

For physical sciences in US, most of these research and development efforts are funded by the US Department of Energy (DOE) and National Science Foundation (NSF). Universities, US National Laboratories, and occasionally industrial partners work together on projects awarded with different flavors of government funds managed under different rules. At Fermilab, we manage multiple such collaborative computing projects for university and laboratory consortia. In this paper, I explore important lessons learned from my experience with these projects. Using examples of projects delivering computing infrastructure for the Lattice QCD Collaboration, I explain how the use of federal enterprise architecture may be deployed to run projects effectively. I also describe the lessons learned in the process. Lessons learned from the execution of the above projects are also applicable to other consortia receiving federal government funds.

Index Terms—Collaborative work, Computer system, Project management, Science

I. INTRODUCTION

THE delivery of big science is dependant not only on the talents of individual scientists, but also on state-of-the-art laboratories and computing capabilities. These facilities, often requiring a large amount of money, allow scientists unprecedented opportunities to advance their knowledge and improve the quality of human life. These very large facilities are required to validate scientific theories by conducting very large scale computational simulations or to run large experiments. These mega-projects require large-scale computational facilities and complementary infrastructures of

network and software. In U.S.A, these research and development efforts are funded by US Department of Energy (DOE) and National Science Foundation (NSF). Executed as collaborative projects, among universities and US National Laboratories, the money comes in different flavors of funding.

At Fermilab, we manage multiple national computing projects for university and laboratory consortia. The emphasis for these projects is on large-scale computing for scientific discoveries. In this paper, I explore important lessons learned from our experiences with one of the major DOE Information Technology (IT) projects designated as an Office of Management and Budget (OMB) major IT investment. I also review related DOE projects to provide insight to the evolving nature of the project management and accountability that may impact scientific collaborations of the future. The first important lesson is awareness of the changing culture of governance required by the funding agencies. It is critical to plan for the governance well from the onset of a project, or even before the award. The second lesson is to plan for technical performance measures, to gather both financial and project management data and to analyze them in a cost-effective manner. The last, but not the least, important lesson is that the success of such collaborative projects depends on the recognition of the diverse institutional management styles of independent organizations making up the consortia. This paper also addresses how key conclusions may be used for planning other large projects for consortia, not necessarily just the major investments.

II. MOTIVATION

During the past fifty years in the history of computing, technology has progressed from computers with processing speed of hundred thousands (10^3) of floating point operations per second (flops) to those capable of tera (10^{12}) flops. However, the current challenge is to construct a computer that can process one peta (10^{15}) flops. These supercomputers are particularly suitable for solving complex mathematical problems and are highly suitable for large-scale simulations. The race for supercomputing excellence began earnestly with the completion of the 40-teraflops Japan Earth Simulator system, which, on completion in 2002, was the fastest computer ever built. As of November 2006, DOE National Nuclear Security Administration owns the world's fastest

supercomputer which delivers a peak speed of 360 teraflops. At this rate, scientist believe that one petaflops of peak speed will be achieved by the year 2010. To take advantage of these supercomputers, various scientific groups are collaborating together to design projects suitable for simulations and analysis experiments in multiple branches of science.

Supercomputers require special purpose algorithms and software to take advantage of their super fast speed. To accelerate the process of software development, Office of Science (DOE) sponsored and funded an initiative titled “Scientific Discovery Through Advanced Computing (SciDAC) during 2001. As described in the DOE web site [1], these “projects are aimed at accelerating research in designing new materials, developing future energy sources, studying global climate change, improving environmental cleanup methods and understanding physics from the tiniest particles to the massive explosions of supernovae.” Universities and other academic entities, including DOE’s National Laboratories, formed consortia to conduct research on computer simulations, grid computing, and other innovative applications leveraging existing advanced computing facilities. These collaborative projects with DOE funds lead the development of necessary infrastructure for the petascale computing. In 2006, building upon the success of the SciDAC initiative, DOE announced the SciDac-2 initiative. For the next three to five years, thirty awarded project teams will conduct research in a variety of fields. National Science Foundation (NSF) joined the initiative with additional funding. These projects depend on various supercomputing facilities available at universities and National Laboratories.

One of the outcomes of the quest for petascale computing is the establishment of the dedicated Lattice Quantum Chromodynamics computing infrastructure facility. This project titled “SC Lattice Quantum Chromodynamics Computing (LQCD),” started in October 2005 to be continued for four years. The project deliverables include the construction and operation of a distributed facility to perform simulation experiments for the Lattice QCD Collaboration. The major objective of this project is to build the most cost-effective computer facility to maximize Lattice QCD simulation products and their analysis. Three National Laboratories formed a consortium to deploy and operate this dedicated special purpose computing facility. This facility consists of computing clusters at Thomas Jefferson National Accelerator Facility (TJNAF) and Fermi National Accelerator Laboratory (FNAL) and the QCDOC computer, a special purpose computer based on IBM’s Blue Gene/L supercomputer at Brookhaven National Laboratory (BNL). This facility is designed to make innovative use of distributed computing resources optimizing parallel programming capabilities needed for Lattice QCD scientists and the phased purchases of modern computing hardware. This project is entirely funded by DOE as a major OBM300 Information Technology (IT) investment Information Technology (IT)

project. The sole customer to this facility is the Lattice QCD Collaboration, made up of high-energy physicists, mostly from universities.

The success of the Lattice QCD facility also depends on the success of the associated SciDAC projects. Currently, the main SciDac-2 project associated with the facilities project is the National Computational Infrastructure for Lattice Gauge Theory project [2]. Funded by the DOE and NSF, it is a separate consortium of universities and National Laboratories. This consortium, responsible for developing algorithms and software, is made up of multiple National Laboratories and universities.

III. COOKING UP THE BIG SCIENCE

To deliver the scientific performance in the collaborative environment described above, it is important to understand the working process of such a joint venture. That is, how and when to process various ingredients to optimize the scientific outcome

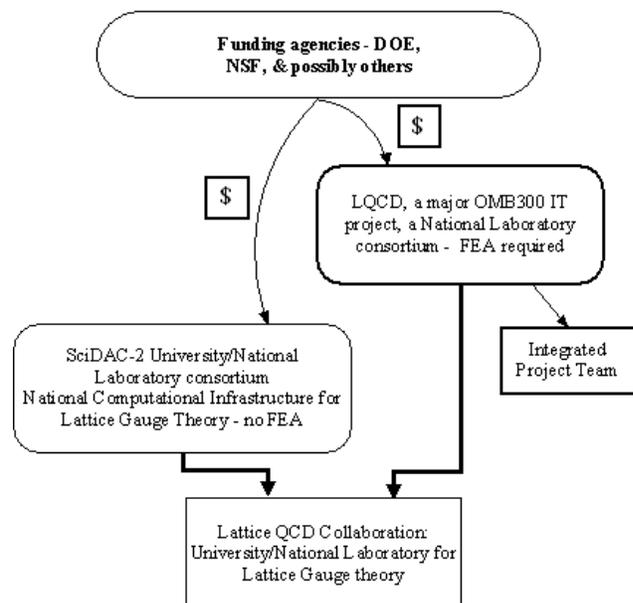


Fig. 1: Components of Scientific Consortia

The above figure shows how various collaborative entities fit together to deliver the scientific results to the Lattice QCD Collaboration. The collaboration is made up of more than twenty-five US universities and five National Laboratories. As shown in the above figure, the SciDAC-2 National Computational Infrastructure for Lattice Gauge Theory project plays a critical role to accomplish the vision of the collaboration. Funded by the DOE and NSF, this project is designated as a non-IT project and does not have to comply with the governance rules of Federal Enterprise Architecture (FEA). The governance required for the SciDAC-2 project

remains minimal. However, the LQCD project executed by three National Laboratories is a major OMB300 IT project and must comply with FEA [3]. The funding level of the SciDAC-2 project and LQCD project is about twenty five million dollars each, making them similar in size. The responsibility for achieving the LQCD performance rests on the Integrated Project Team that includes, among others, the Federal Project Manager, the Contractor Project Manager, and the chair of the Lattice QCD Collaboration. The team assures that all performance objectives, including scientific and financial, are met. Considered to be a medium size project, the number of project components is large enough to learn from them, but small enough that they can be analyzed without adding any additional project management overhead expenses

IV. FEDERAL IT GOVERNANCE AND FEA

The level of governance on major IT projects funded by US government is increasing as the performance accountability expectations are changing. This is true for DOE projects as well [4]. Since the primary focus of scientific projects is to meet the scientific goals, it is necessary to assure that the management processes needed to comply with the governance rules has minimum impact on scientific performance. Laboratories and universities operate under different management principles and structures. Institutional methodologies often differ as well. Establishment of a central project office and a concise, management system for technical and financial performance simplifies the process. The LQCD project office is guided by the Integrated Project Team.

The designation of an IT project as an OMB300 major IT project indicates that the project falls under the jurisdiction of the Clinger-Cohen Act (CCA) [5]. This act is intended to reform the Information Technology Management process and assigns Chief Information Officers of agencies the responsibility to develop and maintain an agency level FEA. This act, initiated by the high degree of failures of government IT projects, was passed to increase customer satisfaction while reducing the cost, complexities, and the lack of interoperability of projects. The objective of the FEA is to implement a uniform IT architecture throughout federal agencies to meet business, technical and financial challenges. As shown in the figure below, the structure of five reference models is at the heart of FEA [3]. FEA guidelines also provide detailed procedures on how to implement the FEA process.

These hierarchically organized reference models of FEA introduced in the Fig. 2 provide management and reporting structure for each IT project. These models help to organize well-defined project parameters. Defining appropriate reference models and refining them is the key to the success of FEA. Under the guidance of the Integrated Project Team of the LQCD project, a set of quantifiable measures is developed in terms of scientific requirements. Related milestones are established and tracked using formal change management

processes.

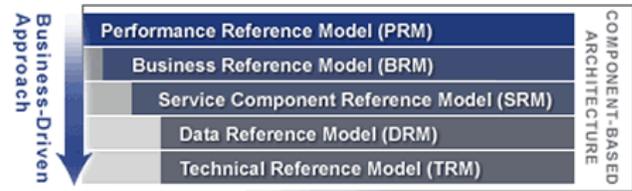


Fig. 2: Federal Enterprise Architecture Reference Model

To obtain federal funding, each OMB300 project is required to submit an annual FEA compatible budget request document, called Exhibit 300, providing the project information based on the reference models. To complete the process, four distinct steps must be completed. These steps are illustrated below using examples from LQCD project:

A) Identify the Line of Business and sub-functions: According to the categories defined by in FEA BRM, LQCD project has three distinct lines of business: Knowledge creation and management; Information and technology management; and General science and innovation.

B) Complete performance reference information (using PRM) table: This information is derived using FEA PRM framework while the measurement groupings are given as sub-functions of BRM. Performance measurement indicators are defined according to the measurement areas and measurement grouping.

The following table shows an example of performance measurement areas and corresponding categories and grouping for the LQCD project.

TABLE 1
PERFORMANCE MEASUREMENT AREAS - EXAMPLES

Measurement Area	Measurement Category	Measurement Grouping
Customer Results	Service Coverage	New Customers and Market Penetration
Mission and Business Results	General Science and Innovation	Scientific and Technological Research and Innovation
Processes and Activities	Cycle Time and Resource Time	Timeliness
Technology	Reliability and Availability	Availability

There are approximately 40 different line items in the LQCD Performance Information Table submitted to OMB. The following table shows examples of measurement indicators for the LQCD project for the fiscal year 2008.

As an example, the measurement indicator “% of generated necessary improved staggered” is a quantification of the simulations to be completed, which is a direct deliverable to

the Lattice QCD Collaboration. Each of these line items has a baseline value, a planned improvement value, and actual results.

TABLE 2
MEASUREMENT INDICATORS – EXAMPLES

Measurement Grouping	Measurement Indicator
New Customers and Market Penetration	Number of distinct users of the facility (includes DOE labs and academic communities)
Scientific and Technological Research and Innovation	% of generated necessary improved staggered – LQCD science deliverables
Scientific and Technological Research and Innovation	Teraflop-Years delivered towards the completion of the 2008 Scientific Program
Timeliness	% of trouble tickets closed within 2 business days
Availability	% of average machine uptime at the clusters and QCDOC

C) Choose service type and components (using SRM): The FEA SRM framework is designed to identify services provided by the investment. Identification of service types and components allow projects to determine services that cover multiple lines of business. LQCD SRMs are functionally independent of the individual sites. This section may also be used to indicate if a service provided by the LQCD project may be reused by another OMB IT project. Service components for LQCD project include library and storage; simulations; and information sharing.

D) Complete technical support information (using TRM): LQCD project places high emphasis on “state-of-the-art” scientific computing dedicated to Lattice QCD simulations. The Table 3 shows a subset of the technical components and their specifications for LQCD.

The culture of formalisms described above is somewhat new to the scientific environment. In the university culture, with floating student population and dependence on grants, it is often difficult to quantify the performance. Establishing a suitable project methodology using the FEA framework was challenging at the beginning of the project. FEA also demands various progress reporting. As time progresses, FEA guides are becoming clearer and accepting the changing culture is becoming easier. The process is beneficial to the centralized project office, as everybody in the consortium can follow the same rule.

The most important management challenge to such projects is the financial one. Annual appropriation of federal funding has an element of uncertainty. Certain projects may be impacted by changes in administrations. Financial challenges must e\be matched with the technical challenge to meet the performance goals. For LQCD project success, efficiency of

the labor force is critical because the labor cost must be minimized to maximize the hardware purchase cost. Consequently, the project management process must be most efficient as well. Technical experts are constantly working on increasing the satisfaction of scientific customers. A small and dedicated group of world-class technical experts continue to fine-tune the facility. They perform benchmarks on machines to maximize the efficiency of the computer hardware purchased and work on improving the robustness of distributed computing facility operation.

TABLE 3
EXAMPLES OF LQCD TECHNICAL SPECIFICATIONS

FEA SRM Component	TRM Service Specification
Simulation	Torque (OpenPBS) Batch System
Information Sharing	Kerberos (MIT krb5)
Data Exchange	International Lattice Data Grid LQCD Metadata Schema
Library / Storage	Jetstore RAID Disk Arrays (SCSI-Attached EIDE and SATA Disk Arrays)
Simulation	Mellanox Infiniband Switches and Host Channel Adapters
Simulation	Custom QCDOC Supercomputer
Simulation	SciDAC Lattice QCD Software Libraries

Considerations for alternate analysis for the investment, including analysis of net present value and life cycle costing are required for all OMB300 projects. The LQCD project uses various computing machine benchmarking and active research on vendor roadmaps to determine alternate analysis. Risk identification and defining mitigation strategies are other mandatory management components. The project, with its stringent investment reporting requirements is, to a degree, a first-of-a-kind project managed at Fermilab. Designing and managing such projects is always challenging [6].

The diversity of people, institutions, and locations imposes another challenge to consortia. Members of Lattice QCD collaboration are located throughout US and operate under different management principles and structures. Out-front efforts to negotiate uniform project management methodology and parameters for the project helped. Solutions to security and privacy issues are becoming increasingly complex and costly for all IT projects and LQCD project is no exception. The project office tracks the compliance to appropriate regulatory issues. The deployment of related tools and procedures is becoming a recurring cost to such consortia.

V. CONCLUSION

The primary lesson is to be cognizant of the changing culture of governance required by the funding agencies. It is critical to take into account the cost of governance appropriately from the onset. It helps in minimizing the governance costs and improves performance of the project.

Secondly, planning and analysis of technical and financial performance measures in a cost-effective manner is crucial. Establishing simple, but concise, tracking processes for technical and financial performance of the consortium as a single entity is necessary.

The last, but not the least, important lesson is to accommodate the diversity of management styles of independent member organizations of the consortium.

ACKNOWLEDGMENT

This work was supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the Department of Energy. The author wishes to acknowledge the support of her supervisor Dr. Stephen Wolbers.

REFERENCES

- [1] Scientific Discovery for Advanced Computing,
<http://www.scidac.gov>
- [2] National Computational Infrastructure for Lattice Gauge Theory;
<http://www.scidac.gov/physics/quarks.html>
- [3] Federal Enterprise Architecture,
<http://www.whitehouse.gov/omb/egov/a-1-fea.html>
- [4] Commission on Engineering and Technical Systems, "Improving Project Management in the Department of Energy (1999)," National Academic Press
- [5] Clinger Cohen Act,
http://www.cio.gov/documents/it_management_reform_act_feb_1996.html
- [6] B. Banerjee, "Estimating Design Costs for First-of-a-kind Projects," Proceedings of IEEE Engineering Management Conference, September, 2006