Mission Need Statement
for the
Next Generation High Performance Production Computing System Project (NERSC-8)

(Non-major acquisition project)

Office of Advanced Scientific Computing Research
Office of Science
U.S. Department of Energy

Date Approved:

____

Month / Year

Submitted by:

____ ____________

David Goodwin, Program Manager — Date
Advanced Scientific Computing Research, Office of Science, DOE

Concurrence:

____ ____________

Daniel Lehman, Director, — Date
Office of Project Assessment, Office of Science, DOE

Approval:

____ ____________

Daniel Hitchcock, Acquisition Executive, Associate Director, — Date
Advanced Scientific Computing Research, Office of Science, DOE

Table of Contents

1 Statement of Mission Need

The U.S. Department of Energy (DOE) Office of Science is the lead federal agency supporting basic and applied research programs that accomplish DOE missions in efficient energy use, reliable energy sources, improved environmental quality, and fundamental understanding of matter and energy. The research and facilities funded by the Office of Science are critical to enhancing U.S. competitiveness and maintaining U.S. leadership in science and technology. One of two principal thrusts within SC is the direct support of the development, construction, and operation of unique, open-access High Performance Computing (HPC) scientific user facilities.

For computing within SC, the Office of Advanced Scientific Computing Research (ASCR) has a mission to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the DOE. Enabling extreme-scale science is a major ASCR priority. In support of this mission ASCR operates the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory. NERSC serves as SC's High Performance Production Computing (HPPC) Facility, supporting the entire spectrum of SC research, and its mission is to accelerate the pace of scientific discovery by providing high performance computing, information, data, communications, and support services for Office of Science sponsored research.

The U.S. Department of Energy (DOE) Office of Science (SC) requires a high performance production computing system in the 2015/2016 timeframe to support the rapidly increasing computational demands of the entire spectrum of DOE SC computational research. The system needs to provide a significant upgrade in computational capabilities, with at least a ten-times increase in sustained performance over the NERSC-6 Hopper system on a set of representative DOE benchmarks.

In addition to increasing the computational capability available to DOE computational scientists, the system also needs to be a platform that will begin to transition DOE scientific applications to more energy-efficient, manycore architectures required for exascale computing. This need is closely aligned with the US Department of Energy's 2011 strategic plan, which states an imperative to continue to advance the frontiers of energy-efficient computing and supercomputing to enable greater computational capacity with lower energy needs. Energy-efficient computing is a cornerstone technology of what has been called exascale computing and represents the only way of continuing NERSC's historic performance growth in response to science needs. In the most recent DOE strategic plan, development and deployment of high-performance computing hardware and software systems through exascale is a targeted outcome within the Science and Engineering Enterprise goal of maintaining a vibrant U.S. effort in science and engineering with clear leadership in strategic areas.

2 Capability Gap/Mission Need

During 2009-2011, ASCR commissioned a series of workshops to characterize the computational resource requirements each program office will need to reach its research objectives in 2014. (For full workshop reports see: http://www.nersc.gov/science/requirements-workshops/.) A careful analysis of these requirements demonstrates a critical mission need for over 14 times the current High Performance Production Computing platform capability by 2014 to address computational needs of projects sponsored by Office of Science Program Offices and to avoid creating an unacceptable gap between needs and available computing resources. By 2016, it is estimated that there will be a need for 47 times the current HPPC capability. Due to budget constraints, NERSC's 2013 supercomputer acquisition ("NERSC-7") will satisfy only a small portion of the need identified in the workshops.

Dr. William F. Brinkman, Director of the Office of Science Department of Energy, stated in March 20, 2012 testimony before the House Subcommittee on Energy & Water Appropriations, that, "In the course of our regular assessment of the needs of the scientific community, it is clear that in several areas DOE's simulation and data analysis needs exceed petascale capabilities." Brinkman, as well as others, noted, however, that making the transition to exascale poses numerous unavoidable scientific, algorithmic, mathematical, software, and technological challenges, all of which result from the need to reduce power consumption of computing hardware.

It is now widely recognized that computing technology is undergoing radical change and that exascale systems will pose a wide variety of technology challenges. These challenges have arisen because hardware chip technologies are reaching physical scaling limits imposed by power constraints. Although transistor density continues to increase, chips will no longer yield faster clock speeds. Instead, vendors are increasing the number of cores on a chip such that within this decade, chips are expected to have over 100 times more processing elements than today. These innovations in architecture and vast increases in chip parallelism (referred to as "manycore") will affect computers of all sizes. Many of the challenges we anticipate in the exascale computing era are confronting us even today. According to the recent National Research Council report *The Future of Computing Performance: Game Over or Next Level?*, "We do not have new software approaches that can exploit the innovative architectures, and so sustaining performance growth—and its attendant benefits—presents a major challenge." Users of such systems will require new algorithms that run more efficiently, exploit parallelism at a much deeper level, accommodate far less memory per process space, and take much greater care in considering data placement and movement and machine resilience. The recent report from the *ASCR Workshop on Extreme-Scale Solvers: Transition to Future Architectures* outlined these challenges in detail for the mathematical solver technology that is at the heart of the application codes that enable DOE scientific discovery.

According to the summary report of the Advanced Scientific Computing Advisory Committee in the fall of 2010 entitled, "The Opportunities and Challenges of Exascale Computing", the transition required to adapt to these new architectures is expected to be as disruptive as that from vector machines to massively parallel systems in the early 1990s. While this transition was not easy, NERSC was able to work with the scientific user community to transform applications from a vector model to an entirely different programming model based on the Message Passing Interface (MPI). A high level of user engagement will again be necessary to help users transition to more energy efficient architectures due to the large number of applications running at NERSC.

Because of NERSC's successful history of engaging with users, locating a production, low-power, manycore machine at NERSC will take DOE scientists through the technology transition to manycore architectures required for exascale. Not all Office of Science computing problems require computing at the exascale level but all SC computing problems must be aware of the all-encompassing issues of energy efficient computing, manycore architectures and programming models, communication reduction, and fault tolerance. In other words, those challenges that have been identified for reaching exascale will be necessary for achieving performance growth on systems of any size but especially at the level required for the next HPC system. Therefore, there is a need to move the entire DOE workload to a more energy-efficient manycore machine; that is, a need to provide, for thousands of users, a machine based on the innovative architectures that are emerging for the next generation extreme scale computing systems. The consequence of the trends noted in the National Research Council report is that computing performance required to achieve DOE science goals can be attained only if parallel computing systems, and the software that runs on them, can become more power-efficient.

2.1 Scientific Demand on Computing Resources

The workshops that NERSC held with each of the SC Offices reinforced the notion that DOE users support projects addressing a wide range of essential scientific issues of pressing national and international concern. The collective nature of the workshops, the selection of participants spanning the full range of Office of Science research areas, and the participation of DOE program managers, enabled derivation of consensus computing resource requirements for science goals that are well aligned with the mission of all six DOE program offices. The workshop reports show that computational requirements for SC Program Offices are expected to dramatically increase and without a significant increase in resources, progress on key DOE SC initiatives are in danger. In total, over 47 times 2011 capability are needed in 2016. This is approximately equivalent to a 75 peak peta-flop/second system. Quite simply, the demand for high performance computing resources by the DOE SC scientists and researchers far outpaces the ability of the DOE Office of Science production facility to provide it.

The scientific goals driving the need for additional computational capability and capacity are clear. Well-established fields that already rely on large-scale simulation, such as fusion and climate modeling, are moving to incorporate additional physical processes and higher resolution. Additional physics to allow more faithful representations of real-world systems is needed, as is the need to model larger systems in more realistic geometries and in finer detail. Acutely constrained resources mean that scientists must compromise between increasing spatial resolution, performing parameter studies, and running higher volumes of simulations.

In order to support the diverse research and development needs of DOE mission-driven science the NERSC-8 system will be expected to support three essential classes of computational inquiry that transcend any one mission category: Science at Scale, Science through Volume, and Science in Data.

2.1.1 Science at Scale

There are many compelling examples within the NERSC workload of science areas for which an extreme scale system is needed. Many of the projects that run at NERSC also have a current INCITE allocation and INCITE users across a broad range of science areas run many of the same codes at NERSC as they do at the Leadership Computing Facilities. There are also many projects at NERSC with codes that run at scale which do not have INCITE allocations. Today, numerous scientists running applications at NERSC have codes capable of routinely using 100,000 cores or more. To illustrate, NERSC workshop attendees reported that requirements for fusion research using particle-based codes will need to use more than a million cores for the largest runs and 10,000 to 100,000 cores for hundreds of routine runs, while integrated transport-

MHD modeling will likely require HPC resources at the exascale to address the problem of electromagnetic turbulence relevant to ITER.

2.1.2 Science Through Volume

A large and significant portion of the scientific discovery of importance to DOE consists of computational science not performed at the largest scales, but rather, performed using a very large number of individual, mutually-independent compute tasks, either for the purpose of screening or to reduce and/or quantify uncertainty in the results.

High-throughput computational screening, in which a database of results from a large number of similar or identical simulations on mixtures of related materials, chemicals, or proteins is compiled and then scanned, has rapidly become a key computational tool at NERSC. One approach is to carry out low-resolution studies on huge numbers of materials to narrow the focus and then carry out more detailed studies on only those identified as promising candidates. This method is being used in the Materials Project to cut in half the typical 18 years required from design to manufacturing, in, for example, the 20,000 potential materials suitable for lithium ion battery construction. Another similar example is the search through approximately three million metal organic frameworks to reveal those most suitable for storing (or sequestering) carbon dioxide; there are currently two successful projects in this area at NERSC. Another example uses massive numbers of molecular dynamics simulations to catalogue protein conformations for better understanding of biological pathways relevant to biofuels, bioremediation, and disease. This method of scientific inquiry is expected to grow at NERSC in the coming years.

A key finding from the NERSC/Basic Energy Sciences (BES) workshop was that poor job turnaround time due to inadequate compute resources profoundly limits the pace of scientific progress in volume-based computational science. The simple availability of additional resources immediately improves the pace of scientific discovery.

Additionally, computational time dedicated to Uncertainty Quantification (UQ) is expected to continue to grow as more scientific disciplines include verification, validation and UQ as standard simulation analysis to provide more defensible and quantifiable understanding of model parameters and simulation results. Already UQ is an important component in simulation areas such as fusion energy, climate modeling, high-energy physics accelerator design, combustion, subsurface environmental modeling, turbulence, and nuclear physics. Many UQ studies require large ensemble calculations. The size of the ensemble runs depends on the number of parameters, model nonlinearities, and the number of models under simultaneous study. In the report High Performance Computing and Storage Requirements for Advanced Scientific Computing Research, experts predict a need for an additional 10 percent increase in computer time across all projects for scientists to perform UQ studies.

2.1.3 Science in Data

A next generation HPC system is required to meet the rapidly growing computational and storage requirements to support key DOE user facilities and experiments such as the Joint Genome Institute (JGI), Planck, and the Advanced Light Source (ALS). As an example, the Office of Basic Energy Sciences (BES) recently created the Scientific User Facilities (SUF) Division, which encompasses 17 user facilities and research centers. In recognition of the growing need for computing resources to support these facilities, BES has created a new pool of computational hours for SUF researchers for the 2012 allocation year. Data set sizes are growing exponentially, because detectors in sensors, telescopes, and sequencers are improving at rates that match Moore's Law. The JGI Microbial Genome and Metagenome Data Processing and Analysis Project is a prime example of the explosive growth in computing needed to support data analysis and also a good example of overlap between Science in Data and Science at Scale. Rapid advances in DNA sequencing platforms driven by comparative analysis and bioenergy needs are fueling a substantial increase in the number and size of genome and metagenome sequence datasets. The next generation HPPC system is required to meet the demand from these facilities.

2.2 Strategic Risk to DOE Office of Science if Not Approved

If the next generation HPPC system is not acquired according to plan and made available to the DOE scientific community on the proposed schedule, there will be numerous, direct, and consequential impacts to the nation's energy assurance and to the DOE mission, such as:

- Approximately 4,500 research scientists will see no increase in computational resources, as required to further scientific discovery, engineering innovation, and maintain U.S. global competitiveness.
- DOE Office of Science HPPC users will not have a system to transition applications and algorithms to the next generation of energy efficient, massive on-chip parallelism systems required for exascale computing.
- DOE Office of Science community will be unprepared to take advantage of energy efficient computing platforms.
- New science initiatives will not be accommodated, thus increasing the number of U.S. scientists and engineers who do not have access to world-class tools, adversely affecting American long-term competitiveness and prestige.
- Other DOE user facilities will be unable to process and analyze data sufficiently or sufficiently quickly and some will be left unanalyzed.

3 Potential Approach

The NERSC-8 project will conduct an Alternatives Analysis prior to CD-1. Potential approaches to meet the mission need include constructing a new facility, utilizing an existing DOE facility such as NERSC, or utilizing a commercial or academic facility. Technical alternatives will include upgrading or replacing an existing computing system, acquiring a new high performance computer system, or utilizing shared resources such as cloud computing. The NERSC-8 project will also consider partnerships with other National Laboratories such as ACES (Alliance for Computing at the Extreme Scale, a collaboration between Los Alamos National Laboratory and Sandia National Laboratory).

NERSC is in a unique position to help transition the broad scientific community to new energy efficient architectures because of its exemplary record of science engagement with users and its support of the broad DOE production computing user base. An extensive 'application readiness' effort will be key to assuring that the architected system can perform well on the DOE workload. The intention of this effort will be to help application developers make changes to their codes that will be sustained to future systems. Doing so is the only way to achieve additional computational performance required to meet science goals.

3.1 Constraints and Limitations

In order to satisfy the mission need, the next generation HPC system must satisfy the following constraints and limitations:

- Maximize application performance while transitioning the DOE Office of Science user community to energy-efficient computing platforms with massive on-chip parallelism.
- Provide high bandwidth access to existing data stored by continuing research projects.
- Delivery of the system in the 2015/2016 timeframe.
- Fit within the budget profile of the NERSC Center.

4 Resource and Schedule Forecast

4.1 Cost Forecast

The estimated Total Project Cost (TPC) range for the considered alternatives is $5 to $19.8M. For this type of project, there is no Other Project Costs (OPC), so TEC is equal to TPC.

4.2 Schedule Forecast

Critical Decision (CD)		Fiscal Year
CD-0	Approve Mission Need	FY 2013
CD-1 CD-3A	Approve Alternate Selection and Cost Range Approve Long Lead-time Procurements	FY 2013
CD-2 CD-3B	Approve Project Baseline Approve Start of Execution (Acquisition)	FY 2014
CD-4	Approve Project Completion	FY 2016

4.3 Funding Forecast

The following funding profile projection is based on the high end of the project cost estimate.

Fiscal Year	FY13	FY14	FY15	FY16	Total ($M)
OPC	0	0	0	0	0
TEC	$2	$3	$8	$6.8	$19.8
TPC	$2	$3	$8	$6.8	$19.8

The Total System Cost is estimated to be between $50M-$100M.