

**ADVANCED SCIENTIFIC COMPUTING
RESEARCH:
DELIVERING COMPUTING FOR THE
FRONTIERS OF SCIENCE**

**FACILITIES DIVISION STRATEGIC PLAN
FOR HIGH PERFORMANCE COMPUTING
RESOURCES**

DRAFT

DRAFT

EXECUTIVE SUMMARY

This plan presents the strategic vision for High Performance Computing (HPC) resources in the Facilities Division of the Advanced Scientific Computing Research (ASCR) program in the Department of Energy's Office of Science for the next 10 years. It responds to the challenge of providing a portfolio of high performance computing (HPC) resources to enable DOE and the Nation's world leadership in critical areas of science, such as the following:

- Developing and optimizing new pathways for renewable energy production and development of long-term secure nuclear energy sources, through computational nanoscience and physics based engineering models.
- Improving our understanding of complex biogeochemical (C, N, P, etc.) cycles that underpin global ecosystems functions and control the sustainability of life on Earth.
- Improving our understanding the roles and functions carried out by microbial life on Earth and adapt these capabilities for human use, through bioinformatics and computational biology.
- Developing integrated modeling environments that couple complex climate models to economic, energy, and resource models that incorporate the human dynamic into large-scale global change analysis.
- Advancing mathematical and algorithmic foundations to support scientific computing in emerging disciplines such as systems biology, molecular self-assembly, emergent behavior of complex systems, agent based modeling and evolutionary and adaptive computing.

ASCR Facilities Division will address these challenges through a balanced program that provides DOE's and the Nation's scientists with high performance production and leadership-class computing resources while fostering the architectural development of the next generation of high end computer hardware and supporting software. The strategy laid out in this "terascale to exascale" plan will focus primarily on HPC computing resources and research and evaluation testbeds as well as the investment needed to stand-up prototypes at the extreme scale –hundreds of petaflops and beyond– for scientific discovery in the 2014-2017 timeframe.

The goal is ambitious, but builds on ASCR's and its predecessors successes over the past fifty years. More recently, HPC capability at ASCR facilities has increased from teraflop computer systems to "hundreds of teraflops" systems. The early HPC resources in our leadership class computing facility are a result of our Advanced Computing Research Testbed program that was instrumental in the testing and maturing of the Cray

X1 product line and of a partnership with the National Nuclear Security Administration (NNSA) on the development of the Red Storm which became the Cray XT3. Through the peer-reviewed Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, DOE is the only agency that has opened its HPC resources at no cost to the entire national and international scientific community including industry to enable breakthrough computational scientific discoveries.

By 2009, scientists will have access to a one petaflop Cray Baker system at the leadership class computing facility at Oak Ridge National Laboratory, a 500 teraflop IBM Blue Gene/P at the leadership class computing facility at Argonne National Laboratory and a 104 teraflop Cray XT4 system at the high performance production facility at Lawrence Berkeley National Laboratory. ASCR's current research and evaluation (R & E) partnership with DARPA on the High Productivity Computer Systems (HPCS) to support the Cray Cascade system and the Blue Gene P/Q development partnership between IBM-Office of Science (SC)-National Nuclear Security Administration (NNSA) are expected to lead to the next generation leadership class resources. In the 2011-2013 timeframe, depending on the availability of funding, either one or both of ASCR's leadership computing facilities will stand-up a "tens of petaflops" HPC resource based on architectures and software developed in these partnerships.

The current major U.S. vendors of large-scale scale systems are in general agreement that extreme scale systems will push the envelope of a number of important technologies, including processor architecture, scale of multicore integration, power management and packaging. Past experience has shown that additional system complexity greatly increases the number of unforeseen hardware and integration challenges associated with installing, managing, and using these machines. Because of the complex nature of these computing systems, during this next 10 years, ASCR must transition to a more holistic view of the facilities. The cornerstone of this vision will be Institutes of Extreme Scale Computing, where pioneering scientific applications and a portion of ASCR's applied mathematics and computer science research are tightly integrated with a portfolio of prototypes, leadership class and production computing facilities. These University-based Institutes represent an evolution of our Research and Evaluation program and will provide training for the next generation of computer scientists, applied mathematicians and computational scientists. To reach our goal, ASCR will embark on an aggressive program to identify the hardware barriers and explore, in partnership with vendors and computer scientists, solutions for overcoming those barriers. More importantly, working with the pioneering application scientists, ASCR will lead the transformation in scientific software and programming techniques to prepare computational scientists to effectively use multicore systems at the 1,000,000 and above processor level for scientific discovery.

The key strategic principles that guide ASCR facility investments include:

- The primary objective of facility investments is to provide support for a broad range of scientific disciplines.

- An integral part of the development of the next generation HPC resources must include a tight coupling with pioneering scientific applications and a portion of ASCR's computer science and applied mathematics portfolio with the development of the next generation of HPC resources .
- A strategic focus to identify and conquer barriers to success while continually evaluating both risk and reward.
- The balance between high performance production and leadership class facilities will be determined with input from the Advanced Scientific Computing Advisory Committee (ASCAC) and the mission needs of the other program offices in the Office of Science

The ASCR program is well-positioned to glue together and make successful government-wide initiatives to revitalize high-end computing in the United States, maintain America's competitiveness and to develop the cyber infrastructure needed to support the national research community. Additionally, ASCR will forge close partnerships with DOE's National Nuclear Security Administration (NNSA) and other Federal agencies—DOD, NSA, DARPA, and NSF—that will build on ASCR's strengths to make unique contributions to the Nation.

Introduction

The mission of the Office of Advanced Scientific Computing Research (ASCR) is to

deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy.

Computational science is increasingly central to progress at the frontiers of almost every scientific discipline and to our most challenging feats of engineering. Computer-based simulation enables scientists to predict the behavior of complex systems that are beyond the reach of our most powerful experimental probes or our most sophisticated theories. Computational modeling has greatly advanced the understanding of fundamental processes of nature, structure, and reactivity. We can now design novel catalysts and high-efficiency engines on computers. Through modeling and simulation, we will be able to explore the interior of stars and learn how protein machines work inside living cells. The other research programs in the Office of Science (SC) depend on the success of ASCR to enable them to answer many of the important questions facing their disciplines. In particular, ASCR's high performance computing facilities provide the "laboratories" for the computational scientists to conduct their studies.

Presidential reports such as the Report of the High-End Computing Revitalization Task Force (HECRTF) and the ASCR advisory committee (ASCAC) on High-Performance Computational Needs and Capabilities in the Office of Science called for substantial investments in high-end computing. In the President's 2005 State of the Union address, supercomputing was identified as an important component in maintaining America's Competitiveness.

Supporting the need for new investments in high-end computing, in 2004 DOE's Office of Science selected the partnership of Oak Ridge National Laboratory, Argonne National Laboratory and Pacific Northwest National Laboratory to establish the first leadership class computing facility at Oak Ridge. Leadership class or capability computing balances ASCR's high performance production or capacity facility at Lawrence Berkeley National Laboratory. By 2009, scientists will have access to a one petaflop Cray Baker system at the leadership class facility at Oak Ridge and a 500 teraflop IBM Blue Gene/P at the leadership class facility at Argonne through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. In the same time period, SC scientists will have access to 104 teraflop Cray XT4 at the high performance production facility at Lawrence Berkeley National Laboratory. It is anticipated that in 2011-2013 ASCR's follow-on "tens of petaflops" HPC resources will evolve from the DARPA HPCS partnership where DOE's investment is supporting the development of the CRAY Cascade system and the IBM-SC-National Nuclear Security Administration (NNSA) Blue Gene development partnership.

ASCR, however, must continue to advance beyond current and near term computational abilities to answer the important scientific questions facing the other

research programs in the Office of Science. To accomplish its mission and to enable critical science to succeed in the extreme scale era, ASCR must address the following challenges:

- Which computational architectures and platforms will deliver the most benefit for the SC scientific strategic goals today and the science of the future as we move toward extreme scale systems?
- What advances in computer science and algorithms are needed to increase the effectiveness of application codes running on multiple cores to solve problems for the Office of Science?
- What operating systems, data management, analysis, model development, and other tools are required to make effective use of future-generation extreme scale supercomputers made up of a large number of possibly heterogeneous processing elements each containing massive multiple core chips?

The solutions to all of these challenges build on the successes of the ASCR program and its predecessors over the past half-century, from the establishment of the applied mathematics research program in the 1950s through the establishment of the first national open supercomputer center in 1974 and the first national leadership class computing facility in 2004. Extrapolation of current hardware trends suggests that the first extreme scale system could be available in the market place by approximately 2022 via a “business as usual” scenario. However, it may be possible to accelerate the availability of these systems for science to approximately 2014-2017 with an aggressive, integrated program, appropriate levels of investments and partnerships with pioneering scientific applications such as climate and fusion energy research.

This following sections details ASCR’s strategy for providing DOE and the Nation with high end computing facilities for open science for the next decade. This discussion includes sections on the ASCR methodology for establishing priorities and allocating facility resources.

ASCR Methodology

The following sections identify important opportunities and ways in which each of the ASCR facilities efforts contribute to the ability of the rest of the Office of Science and the Nation to meet its missions. ASCR ranks and prioritizes these opportunities across the various facilities supporting research efforts to deliver the best program. In addition, because ASCR supports all aspects of the process of scientific discovery, a version of Amdahl's law pertains: "The pace of scientific discovery is determined by the slowest link in the process." Therefore, ASCR investments must also be balanced between the research and facilities. Additionally, the infrastructure required by large-scale facilities such as storage, visualization, data analytics, and networking must be appropriate to the scope of our facilities.

Scope

ASCR recognizes that within the production computing facilities there must exist a range of capabilities, from the computers that can perform the largest and most demanding calculations for the first time anywhere, to the high-performance computers that enable scientists across the Office of Science to do the important science that does not require the most capable computers. ASCR focuses on large-scale facilities that are used by multiple disciplines. We do not consider the smaller mid-range systems that are purchased by individual research groups to be within our mission.

Planning Process

The planning process for ASCR must take into account four important factors:

- The need for input from all of the scientific communities that ASCR facilities support.
- The time it takes to build the infrastructure, i.e. human, physical plant, networking, etc., to support big facilities.
- The special requirements placed on ASCR facilities by opening INCITE to national and international computational scientists, including those in industry.
- The relationship of ASCR facilities to efforts supported by other Federal agencies and other parts of DOE, especially NNSA.

For this reason the ASCR staff must integrate inputs from many sources to develop strategic plans and roadmaps. These sources include Federal staff in other SC program offices; the Advanced Scientific Computing Advisory Committee (ASCAC) and the other Federally chartered advisory committees for SC programs; interagency coordination of information technology R&D under processes established by OSTP; Federal staff at the National Nuclear Security Administration, to ensure close coupling with ASC; and the worldwide scientific research community in scientific disciplines important to DOE.

While the planning horizon for this document is 10 years, it is anticipated that major revisions will be needed more frequently as a result of changes in the external environment, DOE priorities, and the results of evaluations of the ongoing efforts.

Interagency Partnerships

The facilities supported by ASCR are closely coordinated with the information technology research activities of other Federal Agencies (Defense Advanced Research Projects Agency [DARPA], Environmental Protection Agency [EPA], National Aeronautics and Space Administration [NASA], National Institute of Health [NIH], National Security Agency [NSA], and the NSF) through the Networking and Information Technology Research and Development (NITRD) subcommittee of the National Science and Technology Council (NSTC), under the auspices of Office of Science and Technology Policy (OSTP). This coordination is periodically reviewed by the President's Council of Advisors on Science and Technology (PCAST). In addition to this interagency coordination, ASCR has a number of partnerships with other programs in SC and other parts of the Department. Finally, ASCR has a significant ongoing coordination effort with the National Nuclear Security Administration's (NNSA) Advanced Scientific Computing (ASC) Campaign to ensure maximum effectiveness of both computational science research efforts.

One of the results of the HECRTF was a closer coordination of high end computing research, facility and testbed activities. This coordination is required to support a large enough number of research and evaluation prototypes with different architectures (at least 5) to reduce the government-wide mission risk from uncertainty about computer architecture. Coordination is also required to enable scientists to have access to the diverse advanced computers they need to make scientific progress, including systems that enable the largest calculations. These leadership-class systems must be managed by the host agency as a resource for the nation. DOE has established the INCITE program to address this concern.

External Environment

A number of external factors influence the ability of ASCR to achieve its missions.

- The evolution of the commercial market for high-performance computing and networking hardware and software for science.
- Strategic and programmatic decisions made by other (non-DOE) Federal agencies and by international entities.
- The availability of a world-class research community to work on ASCR problems.
- Evolution of public and government attitudes about risk that lead to restrictive security and program evaluation methodologies.
- The fundamentally multidisciplinary, collaborative, distributed, and often international nature of the "big science" questions that DOE is charged with attacking.
- The evolution of government and DOE priorities that affect the level of resources available to accomplish ASCR missions.

The commercial market for high-performance computing and networking hardware and software has a profound effect on the ability of ASCR to meet its missions. We do not fabricate or design any of the underlying technology.

Technology trends and business forces in the U.S. computer system industry over the past decade caused most domestic vendors to curtail or abandon the development of high-end systems designed to meet the most demanding requirements of scientific research. Instead, large numbers of smaller commercial systems have been combined and integrated into large systems that try to achieve the peak performance levels required for agency missions in computational science. The hardware is complicated, unwieldy, and not balanced for scientific applications. Enabling software has been developed for scientists to take advantage of these new computers. However, this software is extraordinarily complex because of the requirements for implementing parallelism. Even if hardware were better optimized for the requirements of science, the inherent complexity of managing large-scale parallelism would make the software complex. Consequently, DOE, primarily through the ASCR, and other Federal agencies whose missions depend on high-performance computing, must make basic-research investments to adapt high-performance computing and networking hardware into tools for scientific discovery.

Investment Strategy

A number of other strategic principles guide ASCR investments:

- Focus on facility investments that support a broad range of scientific disciplines.
- Position pioneering applications as well as a portion of ASCR's computer science and applied mathematics portfolio to be integral parts of the development of the next generation of HPC resources. .
- Employ the principles of Project Management to guide major HPC upgrade project acquisitions.
- A strategic focus to identify and conquer barriers to success while continually evaluating both risk and reward..

To solve next-generation scientific problems, researchers will require a complex portfolio of facilities and testbeds. The elements of this portfolio are listed in Table 1.

Table 1: Elements of Facilities Portfolio

<p>ASCR Facilities</p> <p>High-Performance Production</p> <ul style="list-style-type: none"> • Stable, multiuser environment • Allocations made by SC program offices • Large user support, consulting and training investment • Assist applications preparation to use Leadership Class computing resources • Direct support of agency mission • Resources upgraded every 3-4 years using open competition to vendors • Constant funding stream includes funds for upgrades
<p>Leadership Class</p> <ul style="list-style-type: none"> • Most capable systems available for a class of applications • Small number of projects requiring large, multiyear allocations • Maintain two facilities to provide architectural diversity • Resources for national science community • Allocated through peer-reviewed proposals to INCITE program • Upgrades on 4 year timeframe; dependent on peer-reviewed proposals focusing on scientific opportunity provided by upgrade • 1 or 2 facilities may be upgraded depending on availability of funds
<p>Research and Evaluation Prototypes</p> <ul style="list-style-type: none"> • Peer-reviewed research projects ranging from experimental, proof of concept systems to small prototypes of sufficient scale to enable evaluation of scientific potential • Coordination across DOE (SC/NNSA) and other federal agencies to get needed breadth • 3-5 years in length with no guarantee of follow-on funding • Minimal support for infrastructure • Can be located at laboratories or universities • Close-coupling with computer science and applied mathematics research as well as with application scientists to provide testbeds for developing new algorithms for future architectures. • Enables new science for a few application scientists, who are willing to compute on a research system

In order to address these needs, the current facilities strategy—where procurements of production computers are informed by, but independent of, investments in research and evaluation prototypes or software and algorithmic development—is unlikely to succeed.

We intend to transition to a facilities strategy that requires a much closer coupling of these activities. This strategy will have a more holistic view of the facilities, where a portion of the computer science–network environment research is tightly integrated with a portfolio of experimental, prototype, and production computing facilities. It will also be important to incorporate into this new strategy the lessons learned in our Scientific Discovery through Advanced Computing (SciDAC) and standing up leadership class resources. Primarily, the research and evaluation partnership must include pioneering application scientists who are willing to develop new simulations on experimental systems with erratic operating systems and few software tools. Based on external review, near the end of the evaluation period, some of these research and evaluation systems become candidates either as high performance production systems or as leadership class systems.

It is our vision that ASCR will continue to manage a high-performance production systems to support SC's mission requirements. Leadership-class systems, on the other hand, are resources for the Nation and must be managed as such. The DOE Office of Science believes that a computing capability at this scale is critical for scientific leadership and American Competitiveness. To this end, SC has taken a leadership role in providing this type of facility to the scientific community through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. As with other major facilities, we believe that ASCR should continue to develop and operate leadership-class facilities, with access managed through peer-reviewed proposals submitted to INCITE.

Based on this vision, ASCR has the following strategy for high performance computer facilities and testbeds.

- Continue to evaluate the needs of the application scientists for computing and network facilities. Science is the driver for ASCR facilities. Input will be sought from the broad community of scientists.
- Maintain ASCR's investment in high-performance production computing. This capability is crucial for DOE to accomplish its science mission. Resources at the high-performance production facility are upgraded every 3-4 years through an open competition with vendors. The upgrades are provided for within a relatively constant funding stream.
- Maintain architectural diversity within ASCR's investment in leadership-class computing facilities. This capability is key for DOE and the Nation. Leadership class upgrades are anticipated to occur every 3-4 years and are dependent on peer-reviewed proposals from the current leadership computing facilities focusing on the scientific opportunity provided by the upgrade. Depending on outyear funding profiles, 1 or 2 leadership class resources may be upgraded. It is anticipated that the first solicitation for upgrades will be in the 2008-2009 timeframe.
- Manage the gap between resources available at the high performance production facility and the leadership class facilities to enable the high performance

production facility to succeed in its role in the preparation of applications for leadership class systems.

- Maintain world-class support and consulting organizations for high-performance production and leadership class computing facilities. One of the most critical resources at ASCR facilities is the staff. Effective consulting, user support and training are critical to the scientists working on these facilities. Annual operational assessments of each facility will be used to encourage fiscal responsibility and innovation.
- Increase the number of R&E prototypes under evaluation, and strengthen ASCR's coupling to R&E prototype evaluations funded by other Federal programs such as the DARPA High Productivity Computing Systems (HPCS) program. Another method of increasing R&E prototypes is to support partnerships such as the IBM-SC-NNSA partnership for the follow-on development of the IBM Blue Gene line.
- Manage facility upgrade projects according to the principles of project management. While similar to other Office of Science projects, ASCR facility upgrade projects differ in several ways such
 - "Time" is a key driver throughout the HPCF Upgrade project. The life expectancy of computer technology is measured in years not decades. The resource must be deployed as soon as possible in order to provide the computing capability to the science community before the technology becomes obsolete by advances in the industry.
 - Upgrade projects use innovative financing mechanisms to leverage available funding and to shift the burden of risk from the Government to the vendor (e.g. firm-fixed price, Lease to Own (LTO)). These also represent interaction between the Upgrade project and the Operations Plan to ensure management of financial risk to the entire effort.
 - An HPCF Upgrade project can have significant impact on the steady state operations at the facility, which must be ready for the transition to operations of the upgrade. The reasons for this include the level of integration required and the impact on key staff that play a role in both the operations and the Upgrade project.
- Institute a major review of each R&E prototype at approximately the three-year point to qualify the architecture for installation at either high-performance production or leadership scale. The results of these reviews would be broadly available across the government.
- Improve the coupling between pioneering applications, software research and facilities, with the establishment of Institutes for Extreme Scale Computing.

One critical detail related to this strategy remains to be worked out and that is providing long-term support and maintenance of software. ASCR, as well as a number of other programs in the Office of Science, faces a significant issue in the maintenance of software that results from its research and partnerships. Much of ASCR's software is used by thousands of scientists. The research to develop improved versions of this

software is clearly within ASCR's mission. However, the post development support, maintenance, and testing of software such as PVM, MPICH, PETSC, and a number of linear algebra libraries is a significant challenge. In many respects this software is a new type of virtual facility that must be maintained if the full potential of the software is to be reached.

DRAFT