


Federal Plan for High-End Computing



**Report of the
High-End Computing
Revitalization Task Force
(HECRTF)**



MAY 10, 2004

SECOND PRINTING—JULY 2004

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D.C. 20502

May 10, 2004

Dear Colleague,


The President's FY 2004 Budget articulated what many in the scientific and engineering community have known for some time: high-end computing has become increasingly important to the Nation's scientific and engineering enterprise. High-end computers employing hundreds to thousands of times more computational power than what is available in today's systems are required to solve certain important scientific and engineering problems. These include important national security challenges such as weapon system simulations and image processing of satellite and other data, as well as important scientific and technological questions related to the analysis of complex systems such as aircraft, the atmosphere, and biological systems.

While the importance of high-end computing has grown, the flow of R&D needed to maintain supporting technologies, and the human capital required to sustain them, have not matched this expansion. The current dependence largely on clusters of commercial-off-the-shelf processors and industry's understandable focus on the hardware and software needs of business applications and smaller scale scientific and engineering problems is likely to sustain an existing gap between the Nation's projected high-end computing needs and capabilities.

To begin to address this, the Office of Science and Technology Policy established a task force under the auspices of the National Science and Technology Council's Committee on Technology, made up of Federal agency experts in high-end computing. This High-End Computing Revitalization Task Force was charged with developing a forward-looking plan for high-end computing with the following three components: (1) an interagency R&D roadmap for high-end computing core technologies, (2) a federal high-end computing capacity and accessibility improvement plan, and (3) recommendations relating to federal procurement of high-end computing systems.

The accompanying plan provides these preliminary roadmaps, and initiates a path forward to making high-end computing a vigorous interagency activity that captures the synergies evident in the report—an effort that is already beginning to reap rewards. Addressing the issues facing the Nation's high-end computing enterprise will require a sustained and coordinated effort. The Task Force's report constitutes an important first step.

Sincerely,



John H. Marburger, III
Director

TABLE OF CONTENTS

Executive Summary	iii
1. High-End Computing: A Strategic Tool For Science and Technology Leadership.....	1
The Case for HEC Revitalization	4
Goals.....	6
Scope of the Plan	7
How this Plan Was Prepared	9
HECRTF Charge from OSTP.....	11
2. HEC Research and Development	13
User Requirements for HEC Technology	14
Key Technology Challenges	15
HEC R&D Strategy	17
Hardware Roadmap	18
Software Roadmap.....	19
Systems Roadmap	20
Research and Evaluation Systems	21
Prioritization of HEC R&D Investments.....	21
Agency Participation	24
3. High-End Computing Resources	25
User Requirements for HEC Resources.....	25
Addressing HEC Access, Availability, and Leadership	26
Small Agency Accessibility	26
Large Agency Accessibility	27
Industry Accessibility.....	28
Availability	28
Leadership Systems	29
Recommendations for Access, Availability, and Leadership.....	29
Results of Recommendations for Access, Availability, and Leadership	30
4. HEC Procurement	31
Current HEC Procurement Practice	31
Performance Measures for High-End Systems	32
Total Cost of Ownership	33
Pros and Cons of Consolidating Procurements	34
Procurement Pilot Projects.....	34
Interagency HEC Benchmarking Pilot Project	35
Enhanced Total Cost of Ownership Pilot Project	35
Collaborative Multiagency HEC Procurement Pilot Project.....	35
Procurement Summary.....	36
5. Plan Implementation	37
6. Summary and Conclusions	39

Bibliography.....	41
Appendix A. Sample HEC Applications	43
Appendix A-1. Climate and Weather Research	43
Appendix A-2. Nanoscale Science and Technology	46
Appendix A-3. Life Sciences Applications	49
Appendix A-4. Aerospace Vehicle Design and Optimization	53
Appendix B. Discussion on Time to Solution	57
Appendix C. HEC R&D Technology Areas	59
Appendix D. Discussion on System Specifications and Application Requirements.....	63
Appendix E. An Initial Discussion on Governance of High-End Computing	65
Appendix F. List of Acronyms.....	69
Appendix G. Task Force Membership	72

LIST OF FIGURES

Figure 1: Divergence Problem for HEC Centers	14
Figure 2 HEC Requirements vs. Available Resources for the DoD High Performance Computing Modernization Program	26
Figure 3: Proposed Federal HEC Governance Structure	66

LIST OF TABLES

Table 1-A: Benefits of HEC to Science and Engineering.....	2
Table 1-B: Benefits of HEC to Science and Engineering.....	3
Table 2: Existing Agency Programs	9
Table 3-A: Hardware Roadmap: Current Program	18
Table 3-B: Hardware Roadmap: Robust R&D Plan	18
Table 4-A: Software Roadmap: Current Program.....	19
Table 4-B: Software Roadmap: Robust R&D Plan	19
Table 5-A: Systems Roadmap: Current Program	20
Table 5-B: Systems Roadmap: Robust R&D Plan	20
Table 6: Recommended Priorities	23
Table 7: Impact of Proposed HEC Revitalization Effort	24
Table 8: Agency Participation in HEC R&D	24

EXECUTIVE SUMMARY

Since the World War II era, when scientists, mathematicians, and engineers began using revolutionary electronic machinery that could rapidly perform complex calculations in support of the war effort, pioneering computing capabilities have been a principal foundation of the nation's technological and economic strength.

Today, solving many of our most important scientific and engineering problems requires high-end computers. They are used, for example, for modeling weather and climate patterns; designing highly complex physical systems such as aircraft, ships, and automobiles; and conducting advanced image and signal processing and cryptanalysis. Moreover, U.S. capability in science and engineering is increasingly being called on to address urgent challenges in national and homeland security, economic competitiveness, health care, and environmental protection.

Issues of technology, resources, and governance threaten to limit the potential contributions of high-end computing to vital national interests. In the early 1990s, the Federal government adopted a strategy of pursuing high-end computing capability based on systems built from commercial-off-the-shelf (COTS) components. In the absence of clear evidence against this strategy, the promise of high aggregate performance at relatively low cost made procurement of COTS-based systems a sensible and appropriate course of action. We now have evidence that there are applications of national importance that would benefit significantly from an alternative to COTS-based solutions. Therefore, research and development efforts in alternative architectures and enabling technologies are needed to ensure U.S. leadership in high-end computing (HEC). Further, high procurement and operating costs of HEC systems limit access by many agencies to these resources. A strategy of coordinated acquisition processes coupled with new mechanisms for coordinated access to unique high-end computing facilities is important to ensure the effectiveness and efficiency with which the government acquires these resources.

Recognition is growing in the Administration, the Federal agencies, and Congress that an effective, coordinated national strategy for high-end computing is urgently needed. In 2003, the Department of Defense (DoD), the Department of Energy (DOE), and the National Science Foundation (NSF) initiated independent planning activities to address technology and resource issues. After examining these activities, the Office of Science and Technology Policy (OSTP) determined that an effort focused on high-end computing was warranted. The High-End Computing Revitalization Task Force (HECRTF) was chartered under the National Science and Technology Council (NSTC) to develop a plan for undertaking and sustaining a robust Federal high-end computing program to maintain U.S. leadership in science and technology.

This Plan offers a vision for a proactive Federal effort that advances high-end computing technology to address many of society's most challenging large-scale computational problems and, in doing so, strengthens the nation's global leadership in the sciences, engineering, and technology. The Task Force focused its scope on technology directly needed for high-end computing. Consequently, a number of important technological components essential to science and engineering, such as visualization, networking, grid computing, security, and applications-specific software were not considered by this study. The Plan has three primary components:

HEC Research and Development (R&D), HEC Resources, and Procurement. The basic elements of these three components are summarized below.

HEC Research and Development: The Task Force recommends first and foremost a coordinated, sustained research program over 10-15 years to overcome major technology barriers that limit effective use of high-end computer systems. Today, poor system reliability, the increasing cost and risk of software development, and architectural features (such as the growing imbalance between processor and memory performance) all greatly restrict the ability to achieve high levels of performance for science, engineering, and national security applications. To address these barriers, the HECRTF Plan outlines a comprehensive technology strategy involving basic research, advanced development, engineering and prototype development, and test and evaluation. The Plan presents technology roadmaps for hardware, software, and systems, comparing our current program to a robust revitalization plan. The outcome of the HECRTF R&D Plan will be a robust diversity of tools, technologies, and systems that minimize time to solution for the most challenging computational problems. The end result will be a secure leadership position in high-end computing, and the scientific and technological advances such a position enables, for decades to come.

HEC Resources: Providing high-end computing resources across the full scope of critical Federal missions raises three major issues. First, some agencies have a science and technology mission but lack access to high-end computers. Second, high-end computing has been so successful in contributing to research in science and technology that current resources to meet overall Federal science and engineering demands are heavily oversubscribed. Third, Federal HEC resources do not include systems powerful enough to solve many important large-scale problems. The Task Force provides an interagency collaborative strategy to address all three of these issues.

Procurement: The HECRTF Plan proposes several pilot projects for improving the efficiency of Federal procurement processes, benefiting both government and industry. These pilot projects involve benchmarking (i.e., using software to measure the performance of systems), development of models for total cost of ownership, and approaches to sharing procurement processes across agencies. The intent of the pilot projects is to build teams that span agencies and increase visibility on issues critical to HEC procurement. The Task Force expects that these projects will improve the information flow to assist in the prioritization of future HEC research, development, and engineering investments. Moreover, coordinated procurement of HEC resources will provide more leverage in working with industry vendors to address the needs of the HEC applications communities.

The Plan proposes alternative approaches and planning strategies to carry out these activities. The Task Force analyzes the likely outcomes five years out in the absence of a revitalization effort. The current program allows for some evolutionary advances in high-end computing. However, the Task Force believes that, even with management improvements, the current program will neither maintain U.S. leadership in the face of serious competition nor keep pace with the accelerating growth of demand for high-end computing resources to meet Federal agency needs.

1. HIGH-END COMPUTING: A STRATEGIC TOOL FOR SCIENCE AND TECHNOLOGY LEADERSHIP

In the past decade, computer modeling and simulation of physical phenomena and engineered systems have become widely recognized as the “third pillar” of science and technology – sharing equal billing with theory and experiment. Simulations are performed on computing platforms ranging from simple workstations to very large and powerful systems known as high-end computers.¹ High-end computers enable investigations heretofore impossible, which in turn have enabled scientific and technological advances of vast breadth and depth. High-end computing (HEC) thus has become an indispensable tool for carrying out Federal agency missions in science and technology.

Complex systems such as aircraft, proteins, human organs, nuclear weapons, the atmosphere, and stars can be analyzed and better understood through computer models. With advances in high-end computing power, scientists will be able to model such systems in far greater detail and complexity, and eventually to couple individual models to understand the behavior of an entire system. The opportunity for accelerating progress in many fundamental and applied sciences is compelling.

In view of these opportunities, the Office of Science and Technology Policy (OSTP) determined that an effort focused on high-end computing was warranted. The High-End Computing Revitalization Task Force (HECRTF) was chartered under the National Science and Technology Council (NSTC) to develop a plan for undertaking and sustaining a robust Federal high-end computing program to maintain U.S. leadership in science and technology.

The HECRTF solicited input from leading applications scientists in a variety of disciplines who use high-end computing to advance their research. They were asked to identify important scientific challenges addressable by high-end computing and to estimate the additional computational capability (as a multiple of present high-end capability) needed to achieve the goal.² A summary of the breadth of opportunity for such scientific and technological advances can be found in Tables 1-A and 1-B on pages 2 and 3. The estimates of additional capability needed to achieve the goals ranged from 100 to 1,000 times the current capability of today’s high-end computing resources. Examples of the detailed analysis undertaken by these applications researchers can be found in Appendix A.

For example, fundamental understanding of the emergence of new behaviors and processes in nanomaterials, nanostructures, nanodevices, and nanosystems will require a combination of new theory, new design tools, and high-end computing for large-scale simulation. Similarly, our ability to provide accurate projections of regional climate requires ensembles of simulations on high-end computers at ultra-high resolution, with sophisticated treatment of cloud formation and dispersal, atmospheric chemistry, and regional influences. The intelligence community’s capability to safeguard our nation hinges in part on the ability of high-end computing to tackle diverse computational applications such as cryptanalysis, image processing of satellite and other data, and signal processing for communications traffic, radar, and other signals.

¹ High-end computers are also called supercomputers.

² Computing Research Association, “Workshop on The Roadmap for the Revitalization of High-End Computing,” June 16-18, 2003, and HEC user survey.

Area	Application	Science Challenge	Potential Outcome with 100 to 1,000 Times Current Capability
Physics	Astrophysics	Simulation of astrophysical environments such as stellar interiors and supernovae.	Yield understanding of the conditions leading to the origin of the heavy elements in the universe.
	High-Energy Physics	Achieve a detailed understanding of the effects of strong nuclear interactions so that the validity of the Standard Model can be tested to determine whether physics beyond the Standard Model occurs at extreme sub-nuclear distances.	Guide experiments to identify transition from quantum chromodynamics to quark-gluon plasma.
	Accelerator Physics	Accurate simulations of the performance of particle accelerators.	Optimize the design, technology, and cost of future accelerators, and use existing accelerators more effectively and efficiently.
	Nuclear Physics	Realistic simulations of the characteristics of the quark-gluon plasma.	By developing a quantitative understanding of the behavior of this new phase of nuclear matter, facilitate its experimental discovery in heavy ion collisions.
Nano-science	Catalyst Science/ Nanoscale Science and Technology	Calculations of homogeneous and heterogeneous catalyst models in solution.	Reduce energy costs and emissions associated with chemicals manufacturing and processing. Meet Federally mandated NOx levels in automotive emissions.
	Nanoscale Science and Technology	Simulate the operation of nanoscale electronic devices of modest complexity.	Take miniaturization of electronic devices to a qualitatively new level enabling faster computers, drug delivery systems, and consumer and military electronics.
	Nanoscale Science and Technology	Simulate and predict mechanical and magnetic properties of simple nanostructured materials.	Enable the discovery and design of new advanced materials for a wide variety of applications potentially impacting a wide range of industries.
Aerospace	Simulation of Aerospace Vehicle in Flight	Simulate a full aerospace vehicle mission, such as a full aircraft in maneuver or an RLV in ascent or descent.	Reduce aerospace vehicle development time and improve performance, safety, and reliability.
	Full Liquid Rocket Engine Subsystems Simulation	Simulate full rocket engine subsystems during ascent including turbopump and combustion devices.	Provide capability for risk assessment during Earth-to-orbit and improve safety and reliability of space transportation systems.
	Aviation Systems Simulation	Execute high-fidelity airspace simulations and develop decision system and management tools for terminal area.	Provide capability for effectively managing national airspace and increase safety in terminal area.
Life Sciences	Structural and Systems Biology	Simulations of enzyme catalysis, protein folding, and transport of ions through cell membranes.	Provide ability to discover, design, and test pharmaceuticals for specific targets and to design and produce hydrogen and other energy feedstock more efficiently.
	Signal Transduction Pathways	Develop atomic-level computational models and simulations of complex biomolecules to explain and predict cell signal pathways and their disrupters.	Yield understanding of initiation of cancer and other diseases and their treatments on a molecular level, and the prediction of changes in the ability of microorganisms to influence natural biogeochemical cycles such as carbon cycling and global change.

Table 1-A: Benefits of HEC to Science and Engineering

Area	Application	Science Challenge	Potential Outcome with 100 to 1,000 Times Current Capability
National Security	Signals Intelligence	Model, simulate, and exploit foreign codes, ciphers, and complex communications systems.	Support U.S. policymakers, military commands, and combat forces with information critical to national security, force protection, and combat operations.
	Directed Energy	Advance the directed energy systems design process out of the scientific research realm into the engineering design realm.	Efficiently design next-generation directed energy offensive and defensive weapon systems. Change the design process from years to days.
	Signal & Image Processing & Automatic Target Recognition	Replace electromagnetic scattering field tests of actual targets with numerical simulations of virtual targets.	Design more stealthy aircraft, ships, and ground systems and create the ability to rapidly model new targets, enabling more rapid adaptation of fielded weapon systems' ability to target new enemy weapon systems.
	Integrated Modeling and Test of Weapon Systems	Model complex system interaction in real time with precision.	Replace many expensive, dangerous, and time-consuming ground tests with virtual tests resulting in lower test costs and more rapid development of weapon systems.
Earth and Atmospheric Sciences	Climate Science	Resolve additional physical processes such as ocean eddies, land use patterns, and clouds in climate and weather prediction models.	Provide U.S. policymakers with leading-edge scientific data to support policy decisions. Improve understanding of climate change mechanisms and reduce uncertainty in the projections of climate change.
	Weather and Short-term Climate Prediction	Enable dynamical prediction of frequency and intensity of occurrence of hurricanes/typhoons and severe winter storms 90 days in advance.	Provide critical support to deployed naval, air, and land forces in local, regional, and global combat environments. Lives saved and economic losses avoided due to better severe weather prediction.
	Solid Earth Science	Improved statistical forecasting of earthquake hazards (fault-rupture probabilities and ground motion).	Provide prioritized retrofit strategies. Reduced loss of life and property. Damage mitigation.
	Space Science	Realistically simulate explosive events on the sun, the propagation of the energy and particles released in the event through the interplanetary medium, and their coupling to Earth's magnetosphere, ionosphere, and thermosphere.	Provide decision makers (both civilian and military) with status and accurate predictions of space weather events on time scales of hours to days.
Energy and Environment	Subsurface Contamination Science	Simulate the fate and transport of radionuclides and organic contaminants in the subsurface.	Predict contaminant movement in soils and groundwater and provide a basis for developing innovative technologies to remediate contaminated soils and groundwater.
	Magnetic Fusion Energy	Optimize balance between self-heating of plasma and heat leakage caused by electromagnetic turbulence.	Support U.S. decisions about future international fusion collaborations. Integrated simulations of burning plasma crucial for quantifying prospects for commercial fusion.
	Combustion Science	Understand interactions between combustion and turbulent fluctuations in burning fluid.	Understand detonation dynamics (for example, engine knock) in combustion systems. Solve the "soot" problem in diesel engines.

Table 1-B: Benefits of HEC to Science and Engineering

However, simulations and other computations of this complexity will require hundreds to thousands of times more computational power than is available in today's high-end computers, as well as enhanced software tools and methodologies. The work of developing these systems must begin now.

This Administration recognizes the critical importance of high-end computing. As stated in the "Analytical Perspectives" of the FY 2004 Budget:

Due to its impact on a wide range of federal agency missions ranging from national security and defense to basic science, high-end computing – or supercomputing – capability is becoming increasingly critical. Through the course of 2003, agencies involved in developing or using high-end computing will be engaged in planning activities to guide future investments in this area, coordinated through the NSTC. The activities will include the development of interagency R&D roadmaps for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues (along with recommendations where applicable) relating to federal procurement of high-end computing systems. The knowledge gained from this process will be used to guide future investments in this area. Research and software to support high-end computing will provide a foundation for future federal R&D by improving the effectiveness of core technologies on which next-generation high-end computing systems will rely.

The United States has repeatedly demonstrated that leadership in science and technology is vital to leadership in national defense and national security, economic prosperity, and our overall standard of living. Today, progress in many branches of science and technology is highly dependent on breakthroughs made possible by high-end computing, and it follows that leadership in high-end computing is increasingly crucial to the nation.³

THE CASE FOR HEC REVITALIZATION

There is increasing recognition within the Administration, Congress, and the Federal agencies that development of a viable Federal strategy for revitalizing high-end computing is needed. The importance of high-end computing was stressed by the President's Information Technology Advisory Committee's 1999 report, which asserted as principal findings: "Innovations are required in high-end systems and application-development software, algorithms, programming methods, component technologies, and computer architecture;" and "The high-end computing capability available to the civilian science and engineering community is falling dangerously behind the state of the art." The report also emphasized the need for improvements in sustained performance on real applications, programmability, ease of use, and scalability – the same issues we face today.

As Congressman Sherwood Boehlert, Chairman of the House Science Committee, stated: "[W]e're not at a point of crisis – most of the world's supercomputers are still made by, and used by, Americans. But we are at a pivotal point when we need to make critical decisions to make sure that remains the case."⁴ Several agencies have undertaken studies of the current state

³ The U.S. Congress recognized the importance of high-end computing in the High-Performance Computing Act of 1991 (P.L.102-194), which led to the creation of the Federal High-Performance Computing and Communications (HPCC) Program. The current Networking and Information Technology R&D (NITRD) Program is the successor to the HPCC Program but, as the name implies, has a broader definition.

⁴ "Supercomputing: Is the U.S. on the Right Path?" Full House Science Committee Hearing, July 16, 2003.

of high-end computing and its impact on their missions. These studies resulted in planning activities for investments in high-end computing, both in research and development, and in satisfying mission requirements for access to HEC resources.

Revitalization of high-end computing is needed to refill the research pipeline with new ideas and highly trained people, support the development of robust and innovative systems, and lower industry and end-user risk by undertaking the test and evaluation of HEC systems and software technologies. This revitalization must support advancement across the innovation spectrum – from basic and applied research, to advanced development, to engineering and prototyping, to test and evaluation. Such a comprehensive approach is essential to the establishment of a sustainable research and development process that encourages the generation of competing innovations from the basic research phase, the development of early prototype HEC systems, and the evaluation of these systems on mission-critical test applications.

In July 2002, agencies with a national security mission collaborated to produce the *Report on High Performance Computing for the National Security Community*. The report's development was led by the National Security Agency (NSA) and included the Defense Advanced Research Projects Agency (DARPA), Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP), National Reconnaissance Office (NRO), National Imaging and Mapping Agency (NIMA), Department of Energy/National Nuclear Security Administration (DOE/NNSA), and National Aeronautics and Space Administration (NASA). The report proposed a research, development, and engineering program known as the Integrated High-End Computing (IHEC) Program. The DOE Office of Science also undertook planning activities in calendar year 2002 to assess the need for research and development in HEC software, architecture evaluation, and large-scale computer systems. The capabilities of the Japanese Earth Simulator System, currently the world's most capable high-end computing system, served as a catalyst for the DOE analysis.

These activities have revealed that current high-end computing resources, architectures, and software tools and environments do not meet current needs. Of equal concern, they observe that sustained investigations of alternative high-end systems have largely stopped, curtailing the supply of ideas and experts needed to design and develop future generations of high-end computing systems. High-end computing revitalization will require strong cooperation and collaboration from all quarters with expertise from academic research institutions to the corporate producers of high-end systems needed by the scientific and engineering communities. The analysis and planning from both activities informed the development of this Plan, which the Task Force then extended into a balanced, comprehensive, and coordinated revitalization program including HEC R&D in hardware, software, and systems, and HEC resource procurement, deployment, and management.

The growing importance of high-end computing as a contributor to both national security applications and scientific advancement is well documented by the respective communities. The national security community has consistently specified high-end computing as a critical component of agency missions.^{5, 6, 7} Moreover, various science communities have recently

⁵ Defense Science Board, "Task Force on DoD Supercomputing Needs," Washington, D. C.: Office of the Under Secretary of Defense For Acquisition and Technology, October 11, 2000.

⁶ U.S. Department of Defense (DoD), "Report on High-Performance Computing for the National Security Community," DoD, Washington, DC, July 2002.

⁷ Games, R., "Survey and Analysis of the National Security High-Performance Computing Architectural Requirements," Bedford, MA: MITRE Corporation, June 4, 2001.

specified high-end computing as an important part of their strategic planning. Details, including specific recommendations for strengthening high-end computing in support of these national security and scientific research priorities can be found in several reports in the bibliography.^{8, 9, 10, 11, 12, 13, 14, 15}

GOALS

Scientific and technological advances in many fields are increasingly reliant on large-scale modeling and simulation, and solving many important scientific and technological problems – including some with national security implications – requires a healthy high-end computing environment. In order to revitalize U.S. leadership in high-end computing, the Task Force recommends that the Federal government and its private-sector partners carry out comprehensive, complementary, and synchronized actions over the next five years with the following goals:

- 1) **Make high-end computing easier and more productive to use.** Emphasis should be placed on time to solution, the major metric of value to high-end computing users. Time to solution includes: time to cast the physical problem into algorithms suitable for high-end computing; time to write and debug the computer code that expresses those algorithms; time to optimize the code to the computer platforms being used; time to compute the desired results; time to analyze those results; and time to refine the analysis into improved understanding of the original problem that enables scientific or engineering advances. In addition, a common software environment for scientific computation encompassing desktop to high-end systems will enhance productivity gains by promoting ease of use and manageability of systems. (For a more detailed discussion of time to solution, see Appendix B.)
- 2) **Foster the development and innovation of new generations of high-end computing systems and technologies.** Key research, development, and engineering areas must be nurtured to assure continuous improvement of high-end computing systems that meet the needs of applications. In addition to the traditional research areas of hardware components and systems, the Task Force has identified the requirement for a common system software base to deliver needed improvements in sustained application performance, ease of use, and manageability of high-end systems, as well as a unified software environment for scientific

⁸ National Science and Technology Council, “A 21st Century Frontier for Discovery: The Physics of the Universe,” Office of Science and Technology Policy, 2004.

⁹ National Research Council, *High Energy Density Physics: The X-Games of Contemporary Science*, Washington, D.C.: National Academies Press, 2003.

¹⁰ Nuclear Science Advisory Committee, “Opportunities in Nuclear Science: A Long Range Plan for the Next Decade,” Department of Energy, 2002.

¹¹ National Research Council, Climate Research Committee, *Capacity of U.S. Climate Modeling to Support Climate Change Assessment Activities*, Washington, D.C.: National Academies Press, 1999.

¹² National Research Council, Board on Atmospheric Sciences and Climate, *Improving the Effectiveness of U.S. Climate Modeling*, Washington, D.C.: National Academies Press, 2001.

¹³ Ad Hoc Working Group on Climate Modeling, U.S. Global Change Research Program (USGCRP) “High-End Climate Science: Development of Modeling and Related Computing Capabilities,” Report to the USGCRP, Washington, D.C.: December 2000.

¹⁴ Subcommittee on Global Change Research, Report by the Climate Change Science Program, “Strategic Plan for the Climate Change Science Program Final Report,” July 2003.

¹⁵ Fusion Energy Sciences Advisory Committee (FESAC) Integrated Simulation and Optimization of Fusion Systems (ISOFS) Subcommittee, “Fusion Simulation Project: Integrated Simulation & Optimization of Fusion Systems,” December 1, 2002.

computation encompassing desktop to high-end systems. The Task Force recognizes that these advances will require Federal investment strategies in research and development since the small size of the high-end computing market does not provide sufficient return on investment to the U.S. high-end industrial base.

- 3) **Effectively manage and coordinate Federal high-end computing.** Agencies will be encouraged to plan and operate facilities in a more coordinated fashion, regularly sharing data on the needs of their user communities. The Task Force has identified ways to manage HEC resources that enable agencies to contribute funding and determine usage based on programmatic needs.
- 4) **Make high-end computing readily available to Federal agencies that need it to fulfill their missions.** The cost of high-end computing resources will continue to rise, increasing the need to share the burden across the Federal government, particularly where mission applications and support for user communities overlap. Federal investment in and access to high-end computing must be coordinated across all agencies to maximize the return on investment while minimizing gaps and duplication, and to ensure that various types of high-end computational resources are provided to the broad Federally funded research community in a balanced manner to support each agency's needs.

SCOPE OF THE PLAN

The Plan includes a number of roadmaps outlining all of the core technologies needed for high-end computers that might be manufactured within approximately 15 years. Key elements include:

- Core technology research and development in the hardware, software, and system technologies, including education of new practitioners, that will make high-end computers more powerful, productive, and easy to use
- Capability, capacity, and accessibility strategies to assure that high-end computing resources are readily available to the science and engineering communities that need them
- Efficient procurement strategies that provide high-end computers that meet user requirements

The Plan also describes strategies for improving access to high-end computing resources, including:

- **Leadership Systems:** the leading-edge high-capability computers that will enable breakthrough science and engineering results for a select subset of challenging computational problems. These are problems that have been unsolvable with currently available computing resources.
- **Production Systems:** computers that address the challenging computational problems that require high-end computational resources but do not require access to the extraordinary Leadership Systems

Networking, grid computing, visualization, general security issues, and applications-specific software were considered outside the scope of this planning effort. Procurements of small-scale systems (e.g., a 128-processor cluster used to support local requirements) also were not

included in this planning activity. This enabled the Task Force to focus on HEC technology and acquisition issues. It should be noted that these activities constitute a subset of the agency R&D activities described in the Networking and Information Technology Research and Development (NITRD) Program's Supplement to the President's Budget.¹⁶ The NITRD Program in High-End Computing (as defined in the budget supplement) is funded at approximately \$900 million; however, the activities considered by the HEC Revitalization Task Force represent only about \$158 million of this total. If this effort succeeds, as we believe it will, the revitalization activities discussed in this Plan will have a positive impact on the long-term activities of the entire \$2.6-billion government portfolio for high-end computing.

The Plan proposes alternative approaches and planning strategies for certain activities and discusses the advantages and disadvantages of each approach. The Task Force analyzes the likely outcomes five years out in the absence of a revitalization effort. The current program allows for some evolutionary advances in high-end computing. However, the Task Force believes that, even with management improvements, the current program will neither maintain U.S. leadership in high-end computing in the face of serious competition nor keep pace with the accelerating growth of demand for HEC resources to meet Federal agency needs.

For HEC R&D, the Plan proposes a robust strategy. This strategy offers a path to architectural diversity in 2010 that is responsive to key science and engineering needs and would enable U.S. leadership in all aspects of high-end computing. In addition, the robust strategy outlines a path for addressing the HEC resource capability, capacity, and accessibility needs of the agencies.

The robust strategy also includes a path for acquisition and deployment of HEC resources. These include: (1) Leadership Systems to meet the most extraordinary and demanding computational needs of the Federal government across sectors that include national security, science and technology, and engineering; and (2) Production systems that address computational problems requiring HEC resources below the capability of Leadership Systems. These periodic procurements will capitalize on R&D investments, leverage industry expertise, and help revitalize innovation in high-end computing architectures and technologies for the benefit of the full spectrum of Federal users.

All of the activities described in the HECRTF R&D Plan are focused on the development of dual-use technologies and devices (e.g., computer systems). The intent of the Plan is to support a wide variety of unclassified and classified applications for high-end computing. While some agency requirements dictate that a particular high-end computing system be operated as a classified facility, the underlying technologies developed under the HECRTF R&D Plan will be unclassified and broadly available to academia, government laboratories, and U.S. industry. The only restrictions placed upon technologies and systems developed under this Plan would be covered under the Export Administration Regulations for dual-use goods and articles.

¹⁶ National Science and Technology Council, Committee on Technology, "Networking and Information Technology Research and Development: Advanced Foundations for American Innovation," Supplement to the President's FY 2004 Budget, September 2003.

HOW THIS PLAN WAS PREPARED

The Task Force developed the HEC Revitalization Plan in accordance with the charge prepared by OSTP. Work on the Plan began in March 2003. The Task Force organized itself into four Task Groups:

- Core technology R&D
- Capability, capacity, and accessibility of high-end computing systems
- Procurement issues
- Integration

The Integration Task Group combined the components from the first three Task Groups into a unified whole and prepared the final Plan.

The Task Force included members (listed in Appendix G) from all Federal organizations with a major stake in high-end computing. (The Department of Homeland Security [DHS] is expected to have a significant interest in high-end computing, but during this planning effort the agency was not in a position to be an active participant.) The Plan strengthens the management and coordination of these agencies' programs and research activities to increase the return on existing investments and to maximize the return on proposed future investments. The Task Force examined existing and planned high-end computing programs, as well as other planning activities undertaken by each agency, as summarized in Table 2.

DOE/Office of Science

- Scientific Discovery through Advanced Computing (SciDAC) program
- Advanced Scientific Computing Research (ASCR) program
- Next Generation Architecture (NGA)
- Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program at the National Energy Research Scientific Computing (NERSC) Center
- Calendar year 2002 HEC planning activities

DOE/NNSA

- Relevant portions of the Advanced Simulation and Computing (ASC) program, including the Path Forward and Advanced Architectures initiatives

NSF

- Partnerships for Advanced Computational Infrastructure (PACI) and related research programs
- Cyberinfrastructure Program (HEC Component) and disciplinary research programs in computer architecture, software, and systems*

DoD

- High Performance Computing Modernization Program (HPCMP)
- NSA Research and Development in HEC
- DARPA High-Productivity Computing Systems (HPCS) program
- Proposal for an Integrated High-End Computing Program*

NASA

- HEC-relevant portions of the Computing, Information and Communications Technology (CICT) program

Table 2: Existing Agency Programs (*Planning Activities)

As part of its planning, the Task Force examined similar interagency efforts from the recent past, notably the High-Performance Computing and Communications (HPCC) Initiative and the Next Generation Internet (NGI) Initiative. Several Task Force members had participated in one or both of these initiatives.

The Task Force gathered facts and information from academia, government and national labs, and industry. A total of 84 technical white papers, which provided commentaries from academia, industry, and other non-Federal entities on the HECRTF charge, was solicited. The Task Force participated in an independent workshop, convened by the Computing Research Association on June 16-18, 2003, and attended by more than 200 scientists and engineers from academia, industry, and government. The workshop employed a technical charter similar to the HECRTF charge, consisted of two plenary sessions and eight parallel working group sessions, and established information and prepared recommendations for Task Force consideration.¹⁷ Task Force members also sought information from the high-end computing industry in one-on-one meetings. The Task Force used the information gathered from each of these various forums to ensure the feasibility and practicality of the activities proposed in the Plan.

¹⁷ Computing Research Association, "Workshop on The Roadmap for the Revitalization of High-End Computing," June 16-18, 2003. <http://www.cra.org/reports/supercomputing.pdf>

HECRTF CHARGE FROM OSTP

Coordinated through the National Science and Technology Council, the Task Force is charged with developing a plan to guide future Federal investments in High End Computing (HEC). Based on the needs of important Federally-funded applications of HEC, this plan will lay out an overall strategy for these investments and will include the following areas as coordinated subtasks:

1. High End Computing Core Technologies R&D: This subtask will produce a five-year roadmap, beginning with FY 2005, for core technology development that includes:
 - Identification of key technologies that must be advanced to strengthen the foundation for developing new generations of HEC systems;
 - Coordinated multiagency R&D plans that lay out a set of alternative programs, as well as identification of those agencies that are best suited to carry out each part of the program based on expertise, facilities, or technical priority;
 - Discussion of approaches to planning, selecting participants, and carrying out the research, development, and engineering, in order to enable both revolutionary and evolutionary advances of technology, as well as to enable diffusion of advances in core technologies into commercial industry.
2. Federal High End Computing Capability, Capacity, and Accessibility: This subtask will produce a five-year roadmap, beginning with FY 2005, that includes:
 - Sets of alternative plans for HEC resources that would help to reduce capability or capacity gaps in addressing important applications of HEC;
 - Performance targets for proposed HEC system alternatives that are linked to application domain requirements and user needs;
 - Discussion of the types of system design specifications needed to effectively meet various application domain requirements;
 - Discussion of resources, tools, and techniques needed to minimize “time to solution” by users of HEC systems;
 - Accessibility approaches to make HEC resources available to Federal and non-Federal user communities, as appropriate, beyond the Federal agency that funds or hosts the resources.
3. Federal Procurement of HEC Systems: This subtask will produce findings and recommendations that include:
 - Identification of a strategy for developing practical performance measures for system procurement that correlate well with realized performance of actual applications;
 - Recommended methods for deriving system performance targets from actual or projected application requirements or other user needs;
 - Discussion of total cost of ownership beyond procurement cost, including space, maintenance, utilities, upgradability, etc.;
 - Recommendations for improving processes for acquiring HEC systems based on the above issues.
4. Integration of HEC Strategies: This subtask will produce a five-year roadmap, beginning with FY 2005, for the Federal role in HEC R&D, utilization, and procurement. The roadmap will be based on the needs of important Federally-funded applications of HEC and will include an overall strategy that incorporates appropriate roles for government, academia, and the private sector. This subtask will be closely coordinated with, and based on, the other subtasks.

2. HEC RESEARCH AND DEVELOPMENT

In the early 1990s, Federal support for R&D in HEC technologies stimulated work in high-end computing architectures, hardware, software, and systems. The High Performance Computing and Communications Program supported R&D in both high-end computing and high-speed networking in order to address “grand challenge” scientific and engineering applications.¹⁸

From 1996 to 2001, Federally supported HEC R&D activities waned for two primary reasons: 1) The U.S. government assumed – incorrectly – that industry would fund and conduct most of the research and development required to advance HEC systems and technology. 2) The Federal government did not actively coordinate HEC investments across agencies.

As a result, many promising hardware concepts including superconducting multiprocessors, processor-in-memory (PIM), multithreading, streams, and holographic storage have not yet been incorporated into commercial high-end computing platforms. Progress in HEC software that had occurred has slowed significantly. For example, no follow-on to the parallel programming standards Message Passing Interface (MPI) and OpenMP is foreseen within the next five years. Research in benchmarking and performance modeling tools required to quantitatively assess HEC systems performance and guide future research was delayed for several years.

The high-end computing market is simply not large enough to divert computer industry attention away from the much larger and more lucrative Web-based commerce and business computing sectors. HEC procurements are approximately \$1 billion per year, while the server market by comparison is over \$50 billion per year.¹⁹ If high-end computing is to be revitalized, the Federal government needs to concentrate on research and prototype development to close this gap and provide the advanced technology to meet Federal computing needs.

With industry focused on the lucrative market for servers, and without Federal investments in alternatives, the HEC resources provided by industry have consisted of very large collections of processors designed for smaller systems in the server market. Unfortunately, these massive multiprocessor systems have proven exceptionally difficult to program, and achieving high levels of performance for some important classes of applications has been problematic. Figure 1 on page 14 illustrates this “divergence problem” – the increasing gap between the theoretical peak performance and the sustained system performance (SSP) – for a major high-end computing center.²⁰ Continued technological improvements in microprocessor speeds driven by Moore’s law result in the steeply rising upper curve of theoretical peak performance. However, the result is multiprocessor machines that are increasingly out of balance in terms of processor speed versus memory bandwidth. The imbalance produces the disappointing rise in sustained

¹⁸ *The Federal High Performance Computing Program*, Executive Office of the President, Office of Science and Technology Policy, 8 September, 1989. *Grand Challenges 1993: High Performance Computing and Communications*, Office of Science and Technology Policy, Supplement to the President’s FY1993 Budget.

¹⁹ Kaufmann, Nicholas J., Willard, Christopher G., Joseph, Earl, and Goldfarb, Debra S., “Worldwide High Performance Systems Technical Computing Census,” IDC Report No. 62303, June 2003.

²⁰ The National Energy Research Scientific Computing (NERSC) Center, Lawrence Berkeley National Laboratory. The Sustained System Performance (SSP) is a benchmark designed specifically to reflect the performance of applications codes at the center.

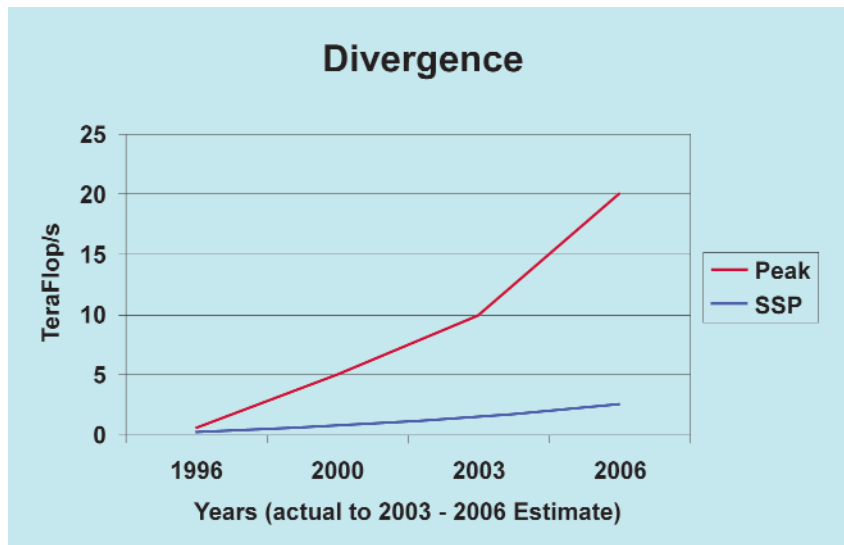


Figure 1: Divergence Problem for HEC Centers (SSP = Sustained System Performance)

system performance displayed by the lower curve. This gap is critical because it is sustained system performance, not peak performance, that is actually usable by applications.²¹

The research and development strategy proposed here offers a new approach that supports the development of robust and innovative systems and lowers industry and end-user risk by supporting the test and evaluation of HEC systems and software technologies. Emphasis is placed on the integration of hardware and software innovations to enable rapid advances in applications.

In one-on-one meetings and at the June 2003 Computer Research Association HECRTF Workshop²², computer industry representatives voiced strong support for U.S. government investments in high-end computing R&D. They argue that investing in breakthrough architectures and technologies and their inclusion in the product development cycle is the most effective approach to influencing systems designs. This strategy helps industry reduce risk, frequently enables vendors to incorporate innovative technologies into their commodity product lines, and creates an improved understanding of industry capabilities and government requirements.

USER REQUIREMENTS FOR HEC TECHNOLOGY

The working group on Application-Driven System Requirements at the HECRTF workshop²³ (consisting of prominent scientists and engineers who conduct scientific research and development using high-end computing) concluded that high-end computing has become essential to advances in many fields. They identified the following primary challenges in effective use of high-end computing:

- Achieving high sustained performance on complex applications
- Building and maintaining complex software applications
- Managing dramatically increasing volumes of data, both input and output
- Integrating multiscale (space and time), multidisciplinary simulations

²¹ This finding has been validated by research funded through DARPA's HPCS program.

²² <http://www.cra.org/Activities/workshops/nitrd/>

²³ *ibid.*

In addition, the working group identified the following goals for future high-end computing systems:

- A 100-fold increase in sustained (as opposed to peak) performance – a level of performance required to solve a number of current scientific and technological problems
- Ultra-fast processors and new algorithms, since not all problems can be easily parallelized
- Improvements in bandwidth and latency for both memory and communications fabric, which for many applications largely determine performance
- A spectrum of architectures to meet diverse application requirements, which can vary dramatically

The lack of tools, programming models, and operating systems software is also a significant concern. The working group concluded that one could expect reasonable performance on up to 1,000 processors with substantial effort, but that programming systems of 100,000 processors (projected in the 2010 timeframe) would be impractical without substantial improvements in software tools, programming models, and algorithms.

Key Technology Challenges

Just as commercial system architectures are driven by the needs of business applications, the design of new generations of HEC systems must be driven by science and engineering applications that support critical mission agency needs and priorities. The Plan focuses on the following three key technology categories and provides a guide to implementation and prioritization in each category:

- **Hardware**, including components and subsystems
- **Software**, including tools and languages
- **Systems**, including architectures and programming models

There are substantial dependencies among these categories that the Plan addresses through a balance of basic and applied research, advanced development, engineering and prototype development, and test and evaluation activities. Brief descriptions are provided below. For a more detailed discussion, please see Appendix C.

Hardware Technology Areas

- **Microarchitecture:** Development of microarchitecture (single chip) designs that better support high-end science and engineering, as opposed to business, workloads. Technical issues include mechanisms for latency hiding, dynamic reconfiguration, and novel processor architectures designed specifically for high-end computing needs.
- **Memory:** Mechanisms for addressing the “memory wall” that is due to the large disparity between the growth of processor speeds (~40% a year) and memory speeds (~7% a year). Technical issues include new cache architectures and intelligent memory controllers.
- **Interconnect:** Improvements in both the bandwidth and latency of the interprocessor links to reduce the performance cost of remote data access. Other issues include intelligent interconnect interfaces and switches.
- **Power, cooling, and packaging:** Methods to reduce the power, cooling, and space

requirements of high-end systems. This will result in smaller systems (i.e., reduced footprint) that deliver higher performance because of reduced communication overhead (signals have to travel shorter distances) and that also have significant life-cycle savings due to reduced power and cooling requirements.

- **I/O and storage:** Methods to address the special I/O and storage requirements of high-end applications including rapid storage and retrieval of extremely large files (terabyte²⁴ to petabyte²⁵ size) through parallel access mechanisms and file systems that are resistant to processing faults.

Software Technology Areas

- **Operating systems:** Development of operating systems that address critical problems in usability, scalability, and reliability of high-end systems. High-end operating systems of the future must be able to scale to hundreds of thousands of processors and enable effective fault-tolerance mechanisms.
- **Languages, compilers, and libraries:** New approaches to writing high-end applications that provide significant improvements in ease of use, interoperability of codes, and portable performance. Compilers that are capable of meeting the demands of high-end applications are sorely needed. Includes development of new languages and automated compiler optimization methods.
- **Software tools and development environments:** Development of new approaches to debugging, performance analysis, and performance optimization that offer significant improvements in ease of use and that are available on all high-end systems. Intelligent development environments that reduce the user's need for detailed, arcane system knowledge and that may also take advantage of application-specific characteristics.
- **Algorithms:** Continued development and improvement of mathematical and computer science algorithms are essential to the success of future generations of high-end architectures. Historically, improved or new algorithms have been a key contributor to performance improvements in science and engineering applications, often rivaling advances in processor speeds.

Systems Technology Areas

- **System architecture:** Development of comprehensive system-wide designs that support high-end science and engineering workloads (rather than solely business workloads). Technical issues include mechanisms for system scaling to 100,000 or more processors, supporting single system images, and ease of programmability.
- **System modeling and performance analysis:** New tools and approaches for analyzing and understanding the interaction between computational requirements of applications and the performance characteristics of proposed new high-end architectures. These tools are needed throughout research and development phases to test and evaluate proposed design alternatives.
- **Reliability, availability, serviceability (RAS), and security:** The large number of processors expected in future generations of high-end systems will pose severe challenges to

²⁴ A terabyte is 10¹² bytes.

²⁵ A petabyte is 10¹⁵ bytes.

reliability, manageability, and security of these systems. Mechanisms are needed to support fault isolation to enable applications to run to completion in the presence of multiple faults, and to provide protection from insider attack and malicious code.

- **Programming models:** Innovative programming models that not only exploit the new capabilities provided by improved architectures but also provide a level of abstraction consistent with the need for significant improvements in programmability. New high-end systems must also effectively support legacy application codes and libraries, frequently based on Fortran, C, and MPI.

HEC R&D STRATEGY

Based on the identified user requirements, the Task Force developed a strategy to sustain a robust technology and industrial base for high-end computing that involves a set of roadmaps detailing the research and development goals and investments. It should be noted that the Task Force explicitly included under R&D the entire cycle: basic research in both innovative HEC architectures and software environments, development of test systems, and engineering of effective scalable HEC architectures. Lessons learned from engineering are then incorporated into new basic research, producing a cyclical innovation process. The current omission of engineering from this process has yielded a broken cycle, impeding the continual improvement of HEC resources to meet the needs of scientists.

Roadmaps and priorities are provided in Tables 3-A through 6 for the hardware, software, and systems technology categories. A more detailed discussion of the basis for these HEC R&D roadmaps is provided in Appendix C. Two R&D planning scenarios are given for each roadmap: the current program, which assumes no resource allocation changes from FY 2004, and a robust component, which assumes a planning, implementation, and procurement process that supports the delivery of new HEC systems in a timely fashion with minimal risk to attaining project goals. It will take a robust investment in HEC R&D to achieve the goals of the Task Force. Both the current program and robust R&D roadmaps provide a projection of high-end computing capabilities in the near term (within a year), the mid-term (within five years), and the long term (within ten years). The near-term capabilities serve as the starting point for future development and are the same for both the current program and robust R&D plans.

The HEC R&D Strategy also includes access to systems for research and evaluation. Such test systems are an important part of the R&D Strategy for three reasons: (1) they provide a testbed, at scale, for early development of new algorithms and computational techniques suited to the proposed architecture; (2) they provide a testbed for software and scalability studies that do not disrupt a production system; and (3) the performance information gained from early evaluation is invaluable for future procurements of full-scale HEC systems.

Finally, software development costs have escalated as users attempt to develop applications codes of ever-increasing complexity using software tools developed a decade or more ago. Improved software tools and technologies are needed across all HEC systems, from clusters to new architectures, to enhance productivity and enable these more complex applications.

Hardware Roadmap

The hardware roadmap (Table 3-A) shows that without additional research effort, as indicated in the current program scenario, there will be little progress beyond the next five years. Improvements in that timeframe depend primarily on industry-driven COTS technology advances and the results of existing or past research investments such as the DARPA Polymorphous Computing Architectures program. Also, there is a strong consensus that, absent significant technological breakthroughs, Moore’s law²⁶ will be coming to an end in the 2015 timeframe.

Current Program		Near-Term	Mid-Term	Long-Term
	Microarchitecture	COTS-driven microarchitecture	Multi-CPU cores per chip, memory bandwidth per CPU decreases	Moore’s law end?
	Interconnect technologies	Interconnect technology based upon electrical interconnect and electrical switches	Interconnect technology based upon electro-optical interconnect and electrical switches	Interconnect technology driven by telecom – expect moderate advances for HEC systems
	Memory	Processor/memory performance gap addressed by caches, limits performance and ease of programming	Early COTS PIM-based and streaming technologies to address processor/memory gap	Evolutionary improvements; increased use of PIMs
	Power, cooling, and packaging	Thermal/packaging – chip/system technologies limited by our ability to cool via air	Evolutionary improvements do not significantly advance our ability to develop high-end systems	System performance limited by “thermal wall”?
	I/O and storage	I/O driven by COTS-based needs in areas of storage and links	Petaflop-scale file systems based upon COTS technologies, RAS issues will limit usability	Depends upon 3-D storage

Table 3-A. Hardware Roadmap: Current Program

Robust R&D Plan		Near- to Mid-Term	Long-Term
	Microarchitecture	Prototype microprocessors developed for HEC systems available	Innovative post-silicon technology optimized for HEC
	Interconnect technologies	Interconnect technology based upon optical interconnect and electrical switches	All-optical interconnect technology for HEC
	Memory	Memory systems developed for HEC needs. Accelerated introduction of PIMs	Revolutionary high-bandwidth memory at petaflop scale
	Power, cooling, and packaging	Stacked 3-D memory and advanced cooling technologies address critical design limitations	Ability to address high-density packaging throughout the entire system
	I/O and storage	Petaflop-scale file systems with RAS focused on HEC requirements	Revolutionary approaches to exascale “file systems”

Table 3-B. Hardware Roadmap: Robust R&D Plan

²⁶ Moore’s Law describes a doubling in capability at fixed cost every 18 months.

Software Roadmap

The current program scenario of the software roadmap (Table 4-A) is dependent upon DARPA's High- Productivity Computing Systems (HPCS) program for the release of new architectures in the next five years that will be more scalable and easier to program. Since the DARPA program ends in 2010, under the current program scenario future progress will be based primarily on those architectures. COTS-based cluster systems will remain a challenge due to minimal investment in the current program.

Current Program		Near-Term	Mid-Term	Long-Term
	Operating systems (OSs)	OSs adapted from desktops or servers. Fragile and do not scale over 1,024 processors	Early introduction of OSs that scale to 10,000 processors for at most two HPCS system architectures. Clusters remain a challenge.	Little progress is expected.
	Languages, compilers, and libraries	Legacy languages and libraries (for example, Fortran, C, C++, and MPI). Compiler technology inadequate for achieving scalable parallelism.	Limited production quality compilers (for example, UPC and Co-Array Fortran [CAF]) for a few systems. MPI continues to dominate. Heroic programming required for computations on over 2,048 processors.	Limited additional improvements in programmability. Production-quality compilers for UPC and CAF widely available. Mostly incremental progress with compiler optimization and MPI implementation. No revolutionary advances in languages
	Software tools and development environments	Wide variety of vendor-specific or research-quality tools – limited integration, difficult to use, and little portability. No integrated development environments (IDEs) available for HEC systems.	Tool capability lags HEC systems (for example, debugging 250,000-processor jobs). IDE support for small-scale (32-processor) systems only.	Gap between tool capabilities and ability to understand large systems widens. IDE support for mid-range shared memory systems
	Algorithms	Efficient parallel algorithms for some problems (for example, dense linear algebra). Others require deep expert knowledge for efficient implementation.	Improved parallel algorithms for unstructured and sparse problems	Additional progress in mapping algorithms onto advanced architectures

Table 4-A. Software Roadmap: Current Program

Robust R&D Plan		Near- to Mid-Term	Long-Term
	Operating systems (OSs)	New research-quality HEC OSs that address scalability and reliability	Production-quality, fault-tolerant, scalable OSs
	Languages, compilers, and libraries	Optimized for ease of development on selected HEC systems. Research-quality implementations of new HEC languages, supporting multiple levels of abstraction for optimization.	High-level algorithm-aware languages and compilers for automated portability across all classes of HEC systems
	Software tools and development environments	Interoperable tools with improved ease of use across a wide range of systems. First research-quality IDEs available for HEC systems.	IDEs that support seamless transition from desktop to largest HEC systems
	Algorithms	New multiscale algorithms suitable for HEC systems. Initial prototypes of architecture-independent parallel computations.	Automatic parallelization of algorithms for irregular and unbalanced scientific problems. Scaling up of parallel algorithms to enable detailed realistic simulations of physical systems.

Table 4-B. Software Roadmap: Robust R&D Plan

Systems Roadmap

As in the software roadmap, the current program scenario of the systems roadmap (Table 5-A) is dependent upon existing research activities (including HPCS) and progress will be difficult after the next five years. The systems area addresses the integration of hardware and software technologies and includes all of the support mechanisms required to create a viable platform. This includes reliability, availability, and serviceability aspects of HEC systems.

		Near-Term	Mid-Term	Long-Term
Current Program	System architecture	COTS-based systems from 10 to 100 Tflops peak (1,000 to 10,000 processors) with server-class operating systems – fragile and hard to program	At most two DARPA HPCS systems capable of sustained petaflops (up to 100,000 processors or more) on selected mission applications	Evolutionary improvements only beyond HPCS systems
	System modeling and performance analysis	System modeling and performance analysis tools developed but ad hoc, incomplete, difficult to use, and not integrated	Accuracy improvements in models/tools for legacy systems and applications for use by experts. Modeling of HPCS systems faces complexity challenges.	Evolutionary improvements toward ease of use and integration with system
	Programming models	Legacy parallel computing models limit ease of programming. Main model is message passing. “Non-heroic” programming practice: MPI at 64 to 256 and OpenMP at 16 to 128.	Minor progress in parallel computing models. “Non-heroic” programming: MPI-2 feasible for 128 to 512 processors and DSM implementations (UPC, CAF, ...) more widespread and available for 64 to 256 processors.	Incomplete implementation and acceptance of shared memory programming models (for example, UPC and CAF)
	Reliability, availability, and serviceability (RAS) + Security	RAS achieved by defensive user actions (for example, checkpoint/restart) and rescheduling	Limited RAS solutions for up to 1,024-processor systems. Partial fault isolation and better profiling of user behavior to prevent inside attack.	RAS solutions for up to 10,000-processor systems. Some improvements in applications security

Table 5-A. Systems Roadmap: Current Program

		Near- to Mid-Term	Long-Term
Robust R&D Plan	System architecture	Three or more systems capable of sustained petaflops (up to 100,000 processors or more) on wider range of applications. Programming much simpler at large scale. Emergence of adaptable self-tuning systems.	High-end systems capable of sustained 10 to 100 petaflops on majority of applications. Programmable by majority of scientists and engineers. Adaptable self-tuning systems commonplace.
	System modeling and performance analysis	Accurate models/tools for HEC systems and applications. Tools and benchmarks provide better understanding of architecture/application interactions.	Models enable analysis and prediction of software behavior. Automated and intelligent performance and analysis tools and benchmarks widely available and easy to use.
	Programming models	Research implementations of novel parallel computing models. “Non-heroic” programming: MPI follow-on for 1,024 processors and robust DSM implementations (UPC, CAF, ...) widespread and available for 1,024 processors.	Parallel computing models that effectively and efficiently match new or planned architectures with applications. Novel parallel computation paradigms foster new architectures and new programming language features.
	Reliability, availability, and serviceability (RAS) + Security	Semi-automatic ability to run through faults. Enhanced prevention of intrusion and insider attack.	Self-awareness: reliability no longer requires user assistance. Systems will have verifiable multilevel secure environments.

Table 5-B. Systems Roadmap: Robust R&D Plan

Research and Evaluation Systems

The HPCC Program in the early 1990s included a strategy to provide HEC systems to universities and research laboratories for experimental use. These “early access” systems enabled testing of early prototypes and provided development platforms for new algorithms and computational techniques. Individual agencies have begun to revive this practice of acquiring systems for research and evaluation, since it is unreasonable to expect future large-scale systems with 10,000 to 100,000 processors to function without proper development and evaluation. The Task Force recommends the procurement of Research and Evaluation Systems as an integral component of the HEC R&D Strategy.

Early access to Research and Evaluation Systems is itself a vital step in the development of usable HEC systems, and it is unrealistic to expect production work to be accomplished on these systems. Often, the hardware and software are not stable because parts of the system may be in development. With such test and evaluation systems, performance metrics used for acquisition of production IT systems are inappropriate. Instead, metrics appropriate to computer science research should be used.

Research and Evaluation Systems are also necessary to support software functionality and scalability studies. Software development testbeds will frequently need to be separate from application testbeds since software development testing often involves “breaking” the hardware and is therefore in conflict with application testing.

The performance information gained from extensive evaluations of Research and Evaluation Systems is invaluable for successful future procurements of production HEC systems. Even if the system is ultimately unsuccessful, in that useful production work is never accomplished, the research project should be considered a success if the test and evaluation effort provides a robust understanding of the system and its failures. Evaluation projects that identify failed approaches save the government from acquiring systems that simply do not perform as expected. They may also suggest more fruitful approaches to remove the sources of failure. It is therefore critical that performance information obtained from testing of Research and Evaluation Systems be shared among all government supercomputer sites and programs.

Prioritization of HEC R&D Investments

There are four major stages of research and development required to sustain a vigorous and healthy high-end computing environment, representing a continuum from basic and applied research, to advanced development, to engineering and prototyping, to testing and evaluation. Each stage in the research and development cycle of an advanced technology such as high-end computing requires a stable pipeline of expertise coming from the university community. In addition, a coherent strategy for the transition from promising prototypes to viable high-end computing technology should be articulated.

Table 6 provides an overview of the HEC R&D investments recommended by the Task Force to address the needs of Federally funded research and development over the next five years in each category:

- **Basic and Applied Research:** Focus on the development of fundamental concepts in high-end computing, creating a continuous pipeline of new ideas and expertise.

- **Advanced Development:** Select and refine innovative technologies and architectures for potential integration into high-end systems.
- **Engineering and Prototype Development:** Perform integration and engineering required to build HEC systems and components.
- **Test and Evaluation:** Conduct testing and evaluation of HEC software and current and new generations of HEC systems at appropriate scale.

Innovations in hardware and systems have a natural transition from research into industrial manufacturing. However, the strategy for software research and development will require a different approach. Major advances in HEC system software have generally occurred only when academia, industry, and government research laboratories have teamed to solve common problems (e.g., the message passing standard, MPI). The high-end software revitalization strategy should include significant government involvement to ensure long-term maintenance of critical HEC software infrastructure components. Consequently, Long-Term Evolution and Support is included in the software component.

The first column of Table 6 lists the three categories of HEC R&D investments – hardware, software, and systems. The second column lists the type of activity for each category – basic and applied research, advanced development, engineering and prototypes, test and evaluation, and long-term evolution and support. The third column shows current program (FY 2004) funding for each category and activity across all involved agencies. These funding amounts are only for activities within the scope of this Plan. Columns four through eight present the total investments recommended by the Task Force in each of the categories and activities for FY 2006 through FY 2010. In these columns, modest funding increments above the FY 2004 level are shown in yellow. Moderate funding increments above the FY 2004 level are shown in lined blue. Robust funding increments above the FY 2004 level are shown in solid blue. Modest funding redirections from the FY 2004 level are shown in hatched pink.

		Current Program*	Increment compared to HEC R&D Current Program				
		FY2004	FY2006	FY2007	FY2008	FY2009	FY2010
Hardware	Basic and Applied Research	(\$ in millions) \$5					
	Advanced Development	\$5					
	Engineering and Prototypes	\$0					
	Test and Evaluation	\$2					
	Total	\$12					
Software	Basic and Applied Research	\$33					
	Advanced Development	\$21					
	Engineering and Prototypes	\$15					
	Test and Evaluation	\$2					
	Long-term Evolution and Support	\$0					
	Total	\$71					
Systems	Basic and Applied Research	\$4					
	Advanced Development	\$40					
	Engineering and Prototypes	\$1					
	Test and Evaluation	\$30					
	Total	\$75					
TOTALS		\$158 ²⁷					

*Assumes no planning changes from FY 2004

 Modest redirection
  Moderate funding increment
 Modest funding increment
  Robust funding increment

Table 6: Recommended Priorities

²⁷ Note: This total represents the aggregate investment across all agencies in High-End Computing as defined in the *Scope of the Plan* section of the report. This does not include many related activities (such as networking, visualization, and application-specific software development) that are outside the scope of this plan.

Table 7 summarizes the mid-term results expected from the R&D revitalization investments in hardware, software, and systems. Hardware and software research investments are expected to provide breakthrough technology otherwise unavailable. Architecture research cycles based on these breakthroughs will produce fully engineered high-end systems. Test and evaluation of these systems will lower program risk and aid transition to production systems. The resultant architectural diversity will aid in matching systems to application needs.

Current Status	Projected Status Without Revitalization Effort	Impact of Revitalization Effort
<ul style="list-style-type: none"> Economic solutions available only for modest problem sizes Limited architecture diversity driven by business needs Virtually no new high-end architectures in hand Major software barriers to effective use of HEC systems 	<ul style="list-style-type: none"> Limited HEC solutions available via DARPA HPCS Incremental progress in cluster software New ideas, people, and advanced development activities are limited Insufficient software support limits applications advances 	<ul style="list-style-type: none"> Multiple new high-end architectures fully engineered and ready for deployment Increased architecture diversity and vendor opportunity Enhanced likelihood of breakthrough technology Additional test and evaluation lowers program risk Basic research program in HEC architectures and related technologies reestablished Innovative research software simplifies use of high-end systems by scientists and engineers.

Table 7: Impact of Proposed HEC Revitalization Effort

Agency Participation

Table 8 shows which agencies are expected to be most active in each of the research categories described within the HECRTF R&D Plan, based on current and past agency activities. Agency contributions and roles will evolve as required to respond to mission interests and research priorities. In Table 8, the DoD listing covers activities in NSA, DARPA, the DoD HPCMP, and the individual Services. The DOE column covers both NNSA and the Office of Science.

		DoD	DOE	NSF	NASA	NOAA
Hardware	Basic and Applied Research	X	X	X		
	Advanced Development	X	X	X		
	Engineering and Prototypes	X				
	Test and Evaluation	X	X			
Software	Basic and Applied Research		X	X		
	Advanced Development		X	X	X	X
	Engineering and Prototypes	X	X		X	X
	Test and Evaluation	X	X		X	
	Long-term Evolution and Support	X	X			
Systems	Basic and Applied Research	X	X	X		
	Advanced Development	X	X	X		
	Engineering and Prototypes	X			X	
	Test and Evaluation	X	X		X	X

Table 8: Agency Participation in HEC R&D

3. HIGH-END COMPUTING RESOURCES

The Task Force defines “HEC resources” as the acquisition, operations, and maintenance of HEC systems that are aligned with the needs of Federal agency mission applications. This definition includes both HEC production systems as well as Leadership-class systems. Although outside the scope of this Plan, the Task Force assumes that agencies will also invest in the broader computing environment, including networking, applications software development, computational science education, general computing and storage systems, and visualization at levels required to support high-end computing as an effective tool.

Several agencies participating in this Plan already have programs that provide HEC resources for mission applications. However, even the largest programs are hard-pressed to meet the needs of their user communities. Other agencies have scant access to resources, and their user communities either do without or glean resources from others. A robust investment in high-end computing by the agencies will accelerate progress in science and engineering and send a strong signal of Federal commitment, generating increased industrial interest and products that reflect science and engineering needs. It should be noted that no civilian agency in the U.S. currently has access to Leadership-class systems to provide true breakthrough capability for important computational problems.

USER REQUIREMENTS FOR HEC RESOURCES

Examining the HEC requirements of a broad range of scientific disciplines across much of the Federal government identifies two classes of resource issues. The first is architectural availability: The high-end marketplace today is not producing machines with the required capabilities to satisfy the most demanding scientific applications. Where there is substantial overlap between commercial computing needs and scientific needs, vendors are supplying products with astounding performance. However, where scientific or defense needs do not overlap substantially with commercial IT, the product space is lacking. In this area, the government could close the gap by providing the requisite investment in research and prototype development that will ensure the appropriate mix of architectures to meet Federal computing needs. (See Appendix D for a discussion of system specifications and applications requirements.)

The second resource issue concerns the ability of agencies to acquire the HEC capacity needed to address Federal government interests. Requests for high-end computing use at Federal science and technology agencies are about three times current capacity²⁸ and are growing by about 80% per year. This demand can be expected to rise further as more scientists begin to exercise production-quality HEC codes and acquire the ability to move from two-dimensional to three-dimensional simulations (see Figure 2 on page 26), and also as more applications become feasible with improved models and more capable systems. The lack of adequate resources is also affecting the national security community, as documented in the July 2002 *Report on High Performance Computing for the National Security Community*.

²⁸ This three-to-one ratio is consistent with the reported demands on other major user facilities, such as neutron sources, light sources, and telescopes.




Figure 2: HEC Requirements vs. Available Resources for the DoD High Performance Computing Modernization Program. Requirements are a measure of user demand, as derived through a formal requirements-analysis process.²⁹

ADDRESSING HEC ACCESS, AVAILABILITY, AND LEADERSHIP

The Task Force proposes separate approaches to address three distinct problems in the accessibility of HEC resources to scientists and engineers with either direct or indirect Federal government support:

- **Accessibility.** Addressing the lack of access to HEC resources: (a) at small agencies with no investments in HEC resources (e.g., NIST), (b) at large agencies with limited or no HEC resources (NIH), and (c) by industry.
- **Availability.** Reducing the gap between requirements and availability for agencies with HEC resources and programs.
- **Leadership.** Providing access to shared Leadership Systems for the largest, most computationally intensive scientific endeavors within the agencies.

Small Agency Accessibility

Barriers and Opportunities: Some Federal agencies have mission-based supercomputing needs but lack access to high-end computer systems. One barrier to entry is the high initial capital investment in supercomputers and associated storage, software, and networks. A second and more important barrier is investment in skilled personnel. Smaller agencies such as NIST often face these hurdles. Agencies with significant high-end computing infrastructure such as DoD, DOE, and NSF have invested in large supercomputer centers that they use for mission applications. Their investments should be, and are, leveraged by other agencies.

²⁹ Office of the Deputy Under Secretary of Defense for Science and Technology, "Department of Defense HPC Modernization Program FY2002 HPC Centers Systems Performance Metrics Report," 2003.

Approach: For small agencies, allow access to existing supercomputing centers and make these available by mutual agreement with other agencies. Small agencies would partner with the larger agencies via multiyear agreements (ideally, four to five years) for utilization of HEC systems, staff expertise, and user training. This approach enables small agencies to leverage existing expertise, lowers the barrier to entry, and facilitates interagency collaboration. To help alleviate the impact of the additional pressure on resources, a cooperative investment strategy should be pursued among the agencies.³⁰

Large Agency Accessibility

Barriers and Opportunities: The scientific drivers behind requirements for supercomputers at NIH are recent, are increasing rapidly, and are less well understood. Although biomedical supercomputing is still in an early stage of development, NIH grantees, taking the individual initiative to apply for computer time, already consume approximately one-third of the cycles provided by the NSF supercomputing centers at Illinois, Pittsburgh, and San Diego. NSF's planned TeraGrid will most likely receive similar use by the biomedical research community.

The largest portion of the current biomedical supercomputing use is for biomolecular modeling and simulation (see Appendix A-3). This class of application will continue to grow rapidly. Within the past year, there has been growing interaction between the practitioners of biomolecular simulation and their counterparts in the fields of computational nanoscience and technology, and materials science, as these workers have become aware of common and complementary problems in molecular physics that are being attacked in each discipline.^{31,32} It is anticipated that the biomolecular physics/nanoscience/materials science nexus will be both a major driver of high-end computing and a major area of interagency interaction. The areas of high-throughput bioinformatics, data integration, and modeling of complex biosystems will grow explosively as software linking high-throughput data gathering with automated analysis and modeling tools becomes widely available. This software is currently in early development with NIH support through new grant programs instituted in the last few years. To realize the potential of imaging and medical informatics will also require increased computational resources. NIH has launched the planning process for the next stage of growth by preparing an eight- to ten-year plan for bioinformatics and computational biology that outlines the software needs and challenges for the next decade of biomedical computing.³³

Approach: In view of the potential impact of NIH programmatic activities on the present supercomputing capacity of the nation, and the prospect of enormously increased impact in the future, the Task Force is encouraged that NIH is beginning to examine seriously its needs for HEC resources. The NIH planning process can serve as a basis for determining an appropriate strategy.

³⁰ National Research Council, *Cooperative Stewardship: Managing the Nation's Multidisciplinary User Facilities for Research with Synchrotron Radiation, Neutrons, and High Magnetic Fields*, Washington D.C.: National Academies Press, 1999.

³¹ International Conference on Computational Nanoscience and Nanotechnology 2004, <http://www.cr.org/ICCN2004>

³² The "NIH Roadmap: Nanomedicine," <http://nihroadmap.nih.gov/nanomedicine/index.asp>

³³ The "NIH Roadmap," <http://nihroadmap.nih.gov/>

Industry Accessibility

Barriers and Opportunities: Access by industry user communities to high-end computing is limited. High-end computing facilities for these users are found in only a few sectors, such as the oil and aircraft industries. The history of government and industry researchers applying computational fluid dynamics to design advanced aircraft illustrates the value of government and private-sector collaboration on important national problems using high-end computing. Computational fluid dynamics has had a profound impact on aircraft design by reducing time-consuming physical tests and decreasing time to market. Other industrial communities with important high-end computing needs are the chemical, semiconductor, and materials sectors, where robust prediction of properties with quantified uncertainties is essential. For the vast majority of applications, particularly those involving mixtures and complex systems (such as drug-protein interactions, polymer nanocomposites, and industrial lubricants), obtaining necessary data through experiments is difficult, time-consuming, and/or expensive.

High-end computing can meet much of this need for robust property predictions by supplementing experimentation and providing data through virtual measurements in a timely manner at lower cost. Virtual experiments using robust predictive models can generate measurements with quantified uncertainties analogous to those obtained by experiment. Experience and resources for these virtual experiments could be gained through industry-government partnerships.

Approach: This Plan proposes the establishment of a pilot activity to enable joint industry/government partnerships for work on HEC problems. Industry access to Federal HEC resources would be governed using policies analogous to those that currently regulate proprietary access to other unique Federal facilities such as neutron sources and x-ray or synchrotron light sources. Industry researchers collaborating with agency researchers would gain experience with computational modeling and the effective use of computational facilities, carrying back new expertise to their institutions. Metrics should be devised and monitored to determine the value of these pilot activities. The intent is to transfer technology and expertise in high-end computing to industry, motivating future industry investment in its own HEC resources. If successful, such programs could increase the adoption and routine use of high-end computing by U.S. industry.

Availability

Barriers and Opportunities: Agencies such as DoD, NSA, DOE/NNSA, DOE/SC, EPA, NASA, NOAA, and NSF all have programs that provide HEC resources to their respective user communities. Most agencies report that these resources are already oversubscribed.

Approach: The HECRTF Plan calls for an increase in the resources available to Federal mission agencies to acquire and operate high-performance computing capabilities. Efficiency improvements due to faster interconnect and memory bandwidth and lower latency – enabled by the R&D portion of this Plan – will further close the capacity gap. While there will always be unmet demand, the goal is to provide resources to satisfy much of the highest-priority, highest-payoff requests.

Leadership Systems

Barriers and Opportunities: U.S. researchers typically have access to the best tools, including world-class high-end computing capabilities. The Task Force has termed these high-capability computers “Leadership Systems.” Leadership Systems would be designed, procured, and administered to provide computing capabilities to scientific researchers years before equivalent systems become commonplace. They would make possible leading-edge science and engineering research for a select set of challenging, high-payoff, and heretofore unsolvable computational problems. Such systems would enable the United States to be “first to market” with important scientific and technological capabilities, ideas, and software, which would later be found on typical high-end systems, then on servers, and finally on desktops. The spinoffs from Leadership Systems should have at least as much impact as the work done at the facilities themselves, as scaled-down versions of the unique machines are brought to market by U.S. computer vendors. Because Leadership Systems are expensive, typically costing in excess of \$100 million per year to procure and operate, their cost is beyond most agencies.

Approach: Build and operate Leadership Systems for the Federal science and technology community. The goal of such systems is to provide computational capability that is at least 100 times greater than what is currently available. A limited set of scientific applications (perhaps 10 per year) would be selected and given substantial access to such systems. Much smaller time allocations could be available to a wider set of applications (perhaps 50 per year) for pilot experiments in preparation for full-scale runs in the future. By their nature, Leadership Systems could be productive for several years, but they would need to be replaced regularly with new Leadership Systems based on scientific need and on technologies emerging from research and development activities.

Recommendations for Access, Availability, and Leadership

Accessibility Recommendations

Agencies whose researchers currently obtain high-end computing resources from other agencies should examine options for formalizing the provision of these resources to their research communities. Agencies with longstanding HEC programs, including DoD, DOE, and NSF, offer models for doing this. One option for resource sharing among agencies is through memoranda of agreement that include cooperative investment agreements³⁴. Each agency should assess and make arrangements to provide for its own resource needs based on mission priorities. Again, agencies with experience in high-end computing can offer models for such determinations and provide expertise to aid in identifying resource needs and management approaches.

Availability Recommendations

Even agencies with longstanding high-end computing programs are only able to provide about one-third of the computing resources required by their funded researchers and engineers. With the growing importance of computational science, modeling, and simulation across all agencies, the future demand for computational resources will place increasing stress on an already

³⁴ National Research Council, *ibid.*

overburdened system.³⁵ Agencies should examine the value of reallocating resources to respond to this opportunity. Agencies may also wish to assess and adjust the relative balance among research and engineering modes (theory, experiment, computation) to ensure optimal allocation of resources.

Leadership Systems Recommendations

Leadership Systems provide leading-edge computational capability for high-priority research problems and guide the development of the next generation of production systems. Because of the high cost for a Leadership-class system, it is expected that not more than one Leadership System will be procured in any year, and such systems would be managed as national resources for all participating agencies. In particular, Leadership-class systems would be operated as an open user facility, with cooperative stewardship practices similar to those used by a large national resource such as the Hubble Space Telescope or Spallation Neutron Source. Access to the system would be governed by a peer review process managed by a council of representatives from the agencies involved. The Task Force estimates that the R&D program recommended in this Plan will provide sufficient technical advances to support periodic procurements, if warranted by scientific need, of new Leadership Systems as new innovative technologies emerge.

The Task Force understands that the deployment of a Leadership System is a major project like the development of a light source, accelerator, or other major scientific instrument. Therefore, the same level of attention to project management, procurement, operation, and maintenance must be paid.

Results of Recommendations for Access, Availability, and Leadership

Each of the HEC resource recommendations presented above confers substantial benefits to mission applications, including:

- **Application of HEC resources to additional important missions.** Increased access at NIH will enable better management of the growth of high-end computing as a complement to experiments and clinical research, reducing the time and cost of research advances. Increased access at NIST will further economic competitiveness through more effective partnerships with industry to develop improved technology.
- **Easing of pressure on resources.** Increased availability will involve replacement of obsolete systems, which will encourage development of additional HEC applications at all agencies. This will provide stimulus to the high-end computing industry to produce more capable systems, drawing on the Plan's R&D components.
- **Providing HEC multiagency, multifaceted, new Leadership Systems.** This will provide agencies with resources of the highest order for leading-edge applications, thereby securing unquestioned HEC leadership for U.S. science and technology. These systems will also provide a development path for the HEC industry to produce scaled-up versions of successful Research and Evaluation systems.

³⁵ See Figure 2 on page 26 for a representative depiction of growth in demand

4. HEC PROCUREMENT

Acquisition – also called procurement – of high-end computers is a complex activity that places significant technical demands on the Federal sector to formulate requirements and evaluate systems. The process is often seen as slow, burdensome, risky, and expensive. The Task Force’s goal is to reverse this perception by developing approaches for procurement of production computing that reduce the burden on both the government and the vendors. The end result: The Federal government receives a system that performs as expected while providing the best value to the taxpayer.

The emergence of computation as a third approach to research and engineering – along with theory and experimentation – and the maturation of simulation as a means of planning and executing agency missions have led to ever-broader application of HEC techniques across the range of Federal missions. Also, some missions do not require full-time use of HEC systems but have important applications needing periodic access to HEC resources. Agencies’ HEC requirements thus are enormously diverse, and HEC centers have mission-related computational problems to solve that require particular system capabilities. Typically, no one vendor manufactures a complete HEC system including computer, storage, networking, workstations, training, system services, software, and development environment. Rather, the typical HEC system is an integrated collection of components from a variety of vendors, all with various features and capabilities. The acquisition process must recognize this diversity in order to target the correct HEC capability for each situation. (The systems the Task Force has called Research and Evaluation prototypes do not fit the model of HEC procurements and must be evaluated as computer science research projects.)

Government agencies constitute the sole market for the highest-capability machines. Agencies either acquire them directly or support acquisitions through the university and national laboratory communities. For this reason, if no other, the Federal government has a vital interest in maintaining a robust, competitive high-end computing industry.

CURRENT HEC PROCUREMENT PRACTICE

Each agency has its own procurement process but there are consistent themes across agencies, such as the use of benchmarking and total cost calculations. All agencies rely heavily on benchmarking, typically drawing on actual agency applications to predict the performance of a proposed HEC system on the specific agency codes. Single-purpose benchmarks, such as LINPACK, may be included in a broad benchmark suite but not as a critical performance measure. “Peak performance” is often cited in vendor advertising and press releases but should not be used as a significant factor in system acquisition.

All agencies use total cost as a prime factor in acquisition decisions. Total cost includes hardware, software, software maintenance, site preparation, operations, and system administration personnel. It may also include training, networking, storage, and the associated development environment, including workstations.

However, these traditional elements of cost overlook the increasing importance of time to solution as a cost driver over the whole life cycle of a HEC system. The labor-intensive programming challenges of ever-more-complex HEC systems increasingly shape their overall cost. Since the price/performance ratio (total cost vs. benchmark performance) is often the deciding factor among competing systems, efforts to improve and streamline the procurement process must better account for time to solution in the measurement of total cost and benchmark performance.

At the same time, there are real differences among agencies' acquisition processes. Some require a live test demonstration (LTD) in addition to benchmarks. The LTD is more common where a narrow set of codes is used in a mission capacity, as is the case with NOAA's National Weather Service systems. Some agencies use benchmarks that are partially or wholly made up of synthetic codes (i.e., short code segments taken from larger applications codes) rather than the actual application codes. This is particularly true where the actual codes are classified. Because such agency-specific differences are mission-driven, they must be accounted for in streamlined procurements to assure the HEC systems' maximum value and productivity to the government.

Agencies use a variety of contracting approaches, including outright ownership, lease to purchase (LTOP), and simple lease. Where ownership rights are acquired, as in a purchase or LTOP, the full capital cost may be required to be budgeted up-front by the agency. This is often a significant hurdle, leading a growing number of agencies to lease machines or to outsource the high-end computing function and simply buy the service on an annual basis where possible.

With the rapid evolution of the electronics industry, the life cycle for HEC systems is decreasing. Moore's law describes a doubling in capability at fixed cost every 18 months. With this rapid increase in capability, it is often more cost-effective to upgrade or replace older systems than to maintain them. Ten years ago, it was common for a HEC system to have a service life of more than five years. Now, the average life span is about three to four years, often with upgrades during even that short lifetime. This puts additional pressure on the acquisition process to shorten procurement timelines. The IT workforce too is under pressure to migrate software to new systems, putting a premium on writing code that is portable from the beginning.

PERFORMANCE MEASURES FOR HIGH-END SYSTEMS

Benchmarks are the primary tool for comparing vendor offerings in the high-end computer market. Performance benchmarks are included in every HEC acquisition (with the possible exception of advanced evaluation systems) in the Federal government. Agencies use actual applications in the benchmarks where appropriate. Many HEC centers have similar user communities that often share the same or similar benchmark codes. For example, NOAA, NASA, DoD, and DOE all use climate models in their mission.

Sustained system performance is currently the only acceptable performance criterion to measure in procurement selection decisions. Other performance indicators, such as calculated peak performance or performance on a single simple benchmark (such as LINPACK) may be useful indicators for discussion purposes but should not be used as the basis for acquisition decisions. Benchmark performance on actual applications is the best indicator of a system's performance

in an operational environment. This is particularly true where the HEC system will be used principally for a narrow set of applications (e.g., weather prediction systems). For computing centers that run a broad range of applications, representative applications are often selected.

Synthetic benchmarks are the only choice when actual applications cannot be used (for example, applications that are classified for security reasons). Synthetic benchmarks must be carefully designed to ensure that the results are meaningful. Preliminary research, supported by DoD and DOE, on the effectiveness of synthetic benchmarks suggests that they may be good predictors of actual user code performance. While encouraging, these results are still quite preliminary and much work remains.

TOTAL COST OF OWNERSHIP

Accurately assessing the total cost of ownership (TCO) of high-end computers is important. First, it allows accurate comparisons between competing alternatives during an acquisition. Second, it serves the vital planning function of determining the total resources that will be required to provide the resource to users. Third, it serves as a means of comparing the efficiency with which computing centers deliver service to their customers.

TCO includes all financial aspects to provide the high-end computing service, and is broken into four major cost areas: (1) hardware; (2) systems software; (3) maintenance, including space, utilities, personnel, and extra-center communications (networking); and (4) user productivity. The first two areas of TCO are fairly well understood and typically accounted for in HEC acquisition and administration. However, networking costs usually are evaluated only when there is consideration of center consolidation. User productivity remains poorly understood; it is the subject of ongoing research.

Hardware items include HEC systems, processors, memory, storage, and internal center communications. These items historically have dominated the TCO for HEC systems, and acquisitions have generally represented these costs adequately. Personnel costs include operations, systems administration, and user support. These costs, as a portion of total HEC systems costs, have been rising as systems become increasingly complex. An increasing portion of the systems integration burden has been shifted away from the HEC vendor to the center or the user, particularly in the areas of systems software, software engineering, and management tools. It is expected that these costs will continue to increase. The evaluation of these costs has generally been well represented in the acquisition and management of HEC resources.

Extra-center communication costs include the networking cost to link remote users to the HEC center and the cost to link HEC centers together. The proportion of total HEC costs that these represent is likely to increase for three reasons. First, user data sets, both for input and output, are increasing in size and are likely to be stored in centralized discipline-specific data repositories remote to the user. Second, any consolidation of HEC centers, either within or across departments and agencies, will create more remote users since most HEC centers are co-located with large user populations. Third, as grid computing becomes viable, the burden on the network is expected to increase.

Productivity is the cost related to how efficiently personnel use the HEC system. Under the DARPA High Productivity Computing Systems program and in collaboration with the DOE Office of Science Performance Evaluation Research Center, research is being conducted to lay

the groundwork for the development of metrics and measurement techniques to provide quantitative approaches to evaluating productivity.

TCO is a valuable tool during the formulation of strategy. To allocate sufficient resources to properly acquire, maintain, and operate a HEC system, organizations must be able to estimate total cost with a high degree of accuracy. This requires a complete analysis of the projected utilization of the system, particularly to deal with access and data management.

PROS AND CONS OF CONSOLIDATING PROCUREMENTS

Each year, several agencies run independent HEC procurement processes targeted at their specific needs. In FY 2002, NOAA ran a competitive procurement for an operational weather supercomputer valued at \$224.5 million over nine years, or about \$25 million per year; NSF acquires systems for the research community via grants; and DoD HPCMP, DoD Fleet Numerical, DOE/NNSA, EPA, DOE/SC, and NASA acquire HEC facilities through contracts or cooperative agreements.

Over the last three years, the DoD HPCMP has experienced significant benefits from consolidating the HEC acquisitions for its four largest centers into a single annual process. These benefits include an acquisition process time of approximately nine months (vs. typical cycle times of 12-18 months elsewhere); improved leverage with potential vendors, yielding better prices and more capacity delivered to customers; ability to consolidate in-house talent from geographically dispersed teams to focus on the same acquisition; better benchmarks and evaluation; and an annual acquisition cycle that industry can plan for ahead of time. In addition, the consolidated process has resulted in team building within the DoD community, producing a community perspective versus individual center perspectives.

Recent HEC solicitations show some overlap in selection criteria. Most major procurement decisions are based on application benchmark performance, among other factors. In some instances there is overlap among the applications, although the specific codes may differ. From the Task Force's perspective, opportunities clearly exist for improved efficiencies from sharing and leveraging procurement activities across agencies.

There will, and should always be, some diversity on how to approach the procurement of HEC systems. Forcing organizations to work together can be problematic. Differing views on strategy for and acceptance of risk can make collaboration difficult. In some cases, such as classified high-end computing uses, agencies such as NSA are unable to share most elements of their procurement process with others. Finally, there is a real concern that over-centralization can reduce innovation and enforce procedures that are not tailored to agency missions. The Task Force concludes that any interagency collaboration on HEC procurements should be voluntary and conducted using a spiral development mode (i.e., taking small steps and building on lessons learned).

PROCUREMENT PILOT PROJECTS

The Task Force proposes three linked pilot projects designed to: 1) improve the efficiency of acquisitions, 2) improve collaboration and information sharing among agencies, and 3) improve measurement of the productivity of these new approaches. The first two projects are technical

in nature but strategic in their contribution to improving the definition of value for HEC systems. The third project is sociological in nature and involves team building across agencies in an area of mutual interest.

Interagency HEC Benchmarking Pilot Project

In this pilot project, selected agencies with similar HEC applications will develop a single suite of benchmarks based on their applications. This benchmark suite is used in a pilot collective acquisition cycle (see following section) in which, for those agencies participating, computer manufacturers benchmark only this suite. Each participating agency uses the benchmarking results, suitably weighted for its applications, in lieu of individual agency-specific benchmarks. Other aspects of agency procurements are unchanged. The effectiveness and efficiency of this approach will be evaluated.

If the project is successful, follow-on activity would involve development of a progressively more comprehensive set of high-end computing benchmarks to be used in acquisitions by a wider range of agencies. This process of expansion of the benchmarks, evaluation of the results, and further expansion would continue in order to further refine the benchmarks.

This Plan proposes that agencies continue and expand research in benchmarking science, including synthetic benchmarks for representing actual applications that are not available to manufacturers for various reasons including security.

Enhanced Total Cost of Ownership Pilot Project

In this project, a multiagency team would evaluate all elements of TCO (e.g., acquisition and maintenance, personnel, extra-center communications, and user productivity) across several similar systems. The TCO model would include underrepresented costs such as grid and distributed computing, larger and remotely located data sets, long-haul communications, increased costs of systems software and middleware, and user productivity, including applications software development costs. This evaluation will develop “best practices” for determining TCO. If successful, the pilot will be expanded to evaluate TCO for dissimilar systems in order to test the extensibility of this approach.

Often the price/performance ratio (total cost/benchmark performance) is the deciding factor among competing systems. Streamlining and improving the measurement of both benchmark performance and TCO will give a fuller indication of costs, which will improve Federal procurement processes. In addition, the insight gained through this activity would be fed back into HEC research, development, and engineering efforts, focusing investments on high-payoff activities, guided by quantitative data rather than based upon unsubstantiated anecdotal evidence.

Collaborative Multiagency HEC Procurement Pilot Project

Applying the techniques and lessons learned from the first two projects, participating agencies will develop a common solicitation and use a single suite of benchmarks for procurement. Agencies that have requirements for supercomputing but do not operate a supercomputing center are encouraged to partner and cost-share in this acquisition. Each agency weights the

benchmarks in a way that suits its applications and includes other agency-specific factors such as risk, live test demonstrations, programming environment serviceability, and system balance. Each participating agency will apply the relevant portions of the TCO model outlined above. The final step will be to evaluate the success of the pilot according to metrics such as improved buying power, reduced overall labor costs, total procurement time, and ability to meet the requirements of the participating agencies.

PROCUREMENT SUMMARY

Procurement of high-end computing resources is a technically demanding and highly complex activity. The high cost of such systems requires careful attention to all aspects of procurement, including:

- Performance measurement, or benchmarking
- Accurate assessment of the total cost of ownership (TCO)
- Procurement consolidation where appropriate

Advances in each of these areas will lower the risk involved in the high-cost acquisition of high-end computing resources. For that reason, the Task Force recommends three pilot programs that seek to improve the practice in each of these three important areas. The lessons learned from these pilot studies will lead to more effective procurement practices for high-end computers in the future.

5. PLAN IMPLEMENTATION

Currently, Federal high-end computing activities are coordinated through the High-End Computing (HEC) Coordinating Group, which reports to the Interagency Working Group (IWG) on Information Technology Research and Development (IT R&D) established in 1991 under the NSTC. While the purview of the working group spans HEC R&D, applications, and infrastructure, most attention has been paid to HEC R&D. Coordination has emphasized information exchanges, community-building activities, and collaborative R&D, with little focus on integration issues.

The Coordinating Group has achieved some tactical successes, notably the development of Message Passing Interface (MPI), a standard for writing portable programs. In this activity, academia, industry, and government labs teamed together in the 1990s to address the serious problem that application codes could not be ported easily across high-end computers. Working together, they developed the MPI standard and reference implementation. The computer industry adopted this standard almost universally, thus enabling portability and performance for a large number of applications codes.

Most recent successful interagency collaborations have been ad hoc, resulting from interactions between individual Coordinating Group members. Notable examples include the DOE/SC Performance Evaluation Research Center and the DARPA HPCS program.

There are a number of different management structures that could be applied to Federal high-end computing. Appendix E discusses one such example. The Task Force expects that final determination of the governance model will be based upon discussions among agency principals, and will be dependent upon the nature of the activities to be implemented. As such, the materials in Appendix E is notional and serves only as a starting point for discussions.

6. SUMMARY AND CONCLUSIONS

The HEC Revitalization Task Force was tasked by the NSTC to “engage in planning activities to guide future investments” in high-end computing. After reviewing over 80 white papers, holding workshops, meeting with high-end computing industry representatives, and preparing this report, this task is complete.

The overarching conclusion of the Task Force is that action to revitalize high-end computing in the U.S. is needed now. The Federal government’s historical success in motivating HEC R&D, the oversubscription of current HEC resources, the scarcity of alternative architectures for delivering high performance to applications, and the lack of current incentives for industry to engage in HEC architecture research all argue strongly that the Federal government should move to revitalize HEC R&D. The Task Force has outlined several steps toward the goal of a healthy and vibrant high-end computing environment:

We identified **key research and development application areas and user requirements** that guide directions for establishing roadmaps and investments in hardware, software, and systems technologies. The HECRTF R&D strategy that emerged from this examination provides a cohesive plan for meeting technology goals in the 10- to 15-year timeframe.

We identified **gaps in accessibility** of Federal agencies to HEC capabilities to meet agency missions. We have proposed a plan and investment strategies to meet the needs of agencies without in-house HEC resources, agencies whose HEC resources cannot meet current demands, and agencies needing Leadership Systems to advance important science and technology research agendas.

We presented a more **effective approach to HEC procurement** and discussed the **establishment of an Interagency Program Office to manage and integrate** the technology, resource, and procurement process to assure efficient and effective HEC investments and results across government agencies. The Task Force believes that, even with management improvements, the current program will neither maintain U.S. leadership in the face of serious competition nor keep pace with the accelerating growth of demand for HEC resources to meet Federal agency needs.

The U.S. government must re-establish its role as the key proponent of HEC systems for the nation’s scientific and engineering research and development. The ideas presented in this report will do much to refill the idea pipeline to advance high-end computing as the third pillar of scientific research and to advance national security, economic security, and the well being of Americans.

BIBLIOGRAPHY

Ad Hoc Working Group on Climate Modeling, U.S. Global Change Research Program (USGCRP), “High-End Climate Science: Development of Modeling and Related Computing Capabilities,” Report to the USGCRP, Washington, D.C.: December 2000.

Advisory Committee to the Director, National Institutes of Health, Working Group on Biomedical Computing, “The Biomedical Information Science and Technology Initiative,” June 3, 1999. <http://www.nih.gov/about/director/060399.htm>

Committee on Environment and National Resources, Subcommittee on Global Change Research, United States Global Change Research Program, and ad hoc Working Group on Climate Modeling, “High-End Climate Science: Development of Modeling and Related Computing Capabilities,” December 2000. <http://www.usgcrp.gov/usgcrp/Library/models2001/default.htm>

Computing Research Association, “Workshop on the Roadmap for the Revitalization of High-End Computing,” June 16-18, 2003. <http://www.cra.org/reports/supercomputing.pdf>

Dally, William, *The Last Classical Computer: ISAT Study*, Stanford University, August 2001.

DARPA High Productivity Computing Systems (HPCS), <http://www.darpa.mil/ipto/programs/hpcs/index.htm>

Defense Science Board, “Task Force on DoD Supercomputing Needs,” Washington, D.C.: Office of the Under Secretary of Defense For Acquisition and Technology, October 11, 2000.

Department of Energy, Office of Science, “Ultrascale Simulation for Science.” <http://www.ultrasim.info/>

Department of Energy, Office of Science, “Reasserting U.S. Leadership in Scientific Computation,” August 6, 2002. http://www.ultrasim.info/doe_docs/Occasional-Science-Computation.pdf

Department of Energy, Office of Science, “A Science-Based Case for Large-Scale Simulation,” Volume 1, July 2003. <http://www.pnl.gov/scales>

Department of Energy, National Nuclear Security Agency, “Advanced Simulation and Computing – Program Overview,” July 25, 2002. <http://www.nnsa.doe.gov/asc/home.htm>

Fusion Energy Sciences Advisory Committee (FESAC) Integrated Simulation and Optimization of Fusion Systems (ISOFS) Subcommittee, “Fusion Simulation Project: Integrated Simulation & Optimization of Fusion Systems,” December 1, 2002.

Games, R., “Survey and Analysis of the National Security High-Performance Computing Architectural Requirements,” Bedford, MA: MITRE Corporation, June 4, 2001.

Kaufmann, Nicholas J., Willard, Christopher G., Joseph, Earl, and Goldfarb, Debra S., “Worldwide High Performance Systems Technical Computing Census,” IDC Report No. 62303, June 2003.

Kuck, D. J., *High Performance Computing: Challenges for Future Systems*, New York, NY: Oxford University Press, 1996.

Muzio, P. and Walsh, R., “Total Life Cycle Cost Comparison: Cray X1 and Pentium 4 Cluster,” Cray Users’ Group 2003 Proceedings, April 22, 2003.

National Research Council, Computer Science and Telecommunications Board, *Funding a Revolution: Government Support for Computing Research*, Washington, D.C.: National Academies Press, 1999.

National Research Council, *High Energy Density Physics: The X-Games of Contemporary Science*, Washington, D.C.: National Academies Press, 2003.

National Research Council, Computer Science and Telecommunications Board, *The Future of Supercomputing: An Interim Report*, Washington, D.C.: The National Academies Press, 2003.

National Research Council, Board on Atmospheric Sciences and Climate, *Improving the Effectiveness of U.S. Climate Modeling*, Washington, D.C.: National Academies Press, 2001.

National Research Council, Climate Research Committee, *Capacity of U.S. Climate Modeling to Support Climate Change Assessment Activities*, Washington, D.C.: National Academies Press, 1999.

National Research Council, *Cooperative Stewardship: Managing the Nation's Multidisciplinary User Facilities for Research with Synchrotron Radiation, Neutrons, and High Magnetic Fields*, Washington D.C.: National Academies Press, 1999.

National Science and Technology Council, "A 21st Century Frontier for Discovery: The Physics of the Universe," Office of Science and Technology Policy, 2004.

National Science and Technology Council, Committee on Technology, "Networking and Information Technology Research and Development: Advanced Foundations for American Innovation," Supplement to the President's FY 2004 Budget, September 2003.

National Science Foundation's Advisory Committee for Cyberinfrastructure, "Revolutionizing Science and Engineering through Cyberinfrastructure," February 3, 2003.
http://www.communitytechnology.org/nsf_ci_report/

Nuclear Science Advisory Committee, "Opportunities in Nuclear Science: A Long Range Plan for the Next Decade," Department of Energy, 2002.

Office of the Deputy Under Secretary of Defense for Science and Technology, "Department of Defense HPC Modernization Program FY2002 HPC Centers Systems Performance Metrics Report," 2003.

Post, D. and Kendall, R., "Software Project Management and Quality Engineering Practices for Complex, Coupled Multi-Physics, Massively Parallel Computational Simulations: Lessons Learned from the ASCI," Washington, D.C.: DOE Software Quality Forum 2003, March 25-26, 2003.

Simon, H., McCurdy, C.W., Kramer, W.T.C., Stevens, R., McCoy, M., Seager, M., Zacharia, T., Nichols, J., Bair, R., Studham, S., Camp, W., Leland, R., Morrison, J., and Feiereisen, B., "Creating Science-Driven Computer Architecture: A New Path to Scientific Leadership," October 2002.
<http://www.nersc.gov/news/HECRTF-V4-2003.pdf>

Subcommittee on Global Change Research, Climate Change Science Program, "Strategic Plan for the Climate Change Science Program Final Report," July 2003.

U.S. Department of Defense (DoD), "Report on High-Performance Computing for the National Security Community," Washington, D.C.: July 2002.

U.S. Department of Defense (DoD), "White Paper – DoD Research and Development Agenda for High-Productivity Computing Systems," DoD, June 11, 2001.

APPENDIX A. SAMPLE HEC APPLICATIONS

A substantial number of application areas from a diverse set of scientific disciplines and Federal agencies were surveyed by the HEC Revitalization Task Force. Four classes of applications are summarized here: Climate and Weather, Nanoscale Science and Technology, Life Sciences Applications, and Aerospace Vehicle Design. These illustrative application summaries are included in order to provide a sense of the promise that sustained growth in HEC capability holds and to give a flavor of the role that high-end computing plays in the breadth of fields of interest to the Federal government.

APPENDIX A-1. CLIMATE AND WEATHER RESEARCH

Introduction: Climate and weather represent one of the Grand Challenges for high-end computing. This is due to the enormous range of spatial and temporal scales that characterize the Earth's physical and biological processes. Spatial scales range from the molecular scales of chemical processes in the atmosphere and ocean to the planetary-scale dynamics that determine the Earth's climate. Temporal scales range from the minutes required to forecast a tornado to multi-century model integrations used to investigate climate change. In recent years, weather modeling has broadened to include space weather, expanding the physics modeled (magneto-hydrodynamics and charged particles) and increasing the range of spatiotemporal scales. Coupled numerical models of the atmosphere, ocean, sea ice, land, and space provide the only means for predicting Earth system phenomena across these scales. Large volumes of high-resolution observations are used to initialize and update models. Our ability to bring models and data together with sufficient resolution and thus enhance our ability to understand climate phenomena and to forecast weather is limited today by the state of available high-end computing.

Illustrative Problems:

A) Longer-term climate projections at regional and global scales: In addition to gaining a fundamental understanding of climate systems to reduce uncertainty in long-term climate projections, research into climate dynamics is necessary to develop the capacity to predict:

- The likelihood of extreme climate fluctuations and abrupt climate change
- The impact of different scenarios for anthropogenic forcing on climate and ecosystems
- The likelihood, intensity, and duration of droughts
- The influence of different land-use scenarios on regional climate
- The distribution and intensity of atmospheric dust (both as a health hazard and as a strong influence on climate evolution)
- The impact of emission reduction strategies on regional climate scales
- The viability of different carbon sequestration scenarios
- The frequency of severe weather events such as hurricanes and floods.

These types of predictions can aid policymakers in developing adaptation strategies that minimize economic disruption from climate fluctuations. They further form the necessary basis for marketplace valuation of carbon emission and sequestration activities.

Due primarily to limitations in computer resources, current models have coarse resolution in which physically important processes such as cloud dynamics and ocean mesoscale dynamics are parameterized instead of being resolved. Although the fruit of considerable research, the best parameterizations still fail to do justice to the physics that they are trying to represent and produce results that are at variance with observations. Similarly, low resolution is recognized as a significant problem in the land surface component. It is generally believed that resolving these processes will be the basis of the next generation of climate modeling research.

The following three problems illustrate what could be achieved with additional computing resources for climate science:

- 1) **The movement of heat and chemical constituents within the Earth's oceans.** In the ocean, a large part of this transport resides in ocean eddies. These have dynamical scales ranging down to 30 km or less. Eddy-permitting models at a minimum will require at least a 10-fold increase in horizontal resolution from 100 kilometers (km) to 10 km, with an overall required increase of about 1,000x in computational power to integrate for the same amount of model time as current models.
- 2) **Land use and vegetation type, and the interaction of these with the physical climate system.** Important processes include conversion of forest to agriculture, desertification, and deglaciation. Resolving typical land use and vegetation scales requires a 10-fold increase in horizontal resolution from 100 km to 10 km, and an overall increase of about 1,000x in the computational performance of the land model.
- 3) **Unresolved processes in the Earth's atmosphere.** Physical processes such as convection in the tropics, gravity-wave breaking, atmospheric boundary layer processes, precipitation, and formation of atmospheric aerosols and their effect on radiative properties are heavily parameterized in current climate models. Simulating the effects of cloud processes will require an atmospheric model with 2 to 5 kilometers resolution (20 to 30 times that used in current models), requiring an overall increase of 8,000 to 27,000 times computing power to match current model integration times. This does not consider increases in vertical resolution or the increased computations required for additional physical processes.

B) Seasonal-interannual (S-I) variability and prediction: The goals for this research include improving the skill of climate predictions on the 6- to 12-month time scale. As with models used in decadal-centennial climate change research, current S-I climate prediction systems consist of ocean-atmosphere-land-sea-ice models. However, the shorter period over which predictions are needed generally permits the use of higher-resolution models, and assimilation systems are required because the memory of initial conditions, particularly over land and oceans, is retained over S-I timescales. Assimilation brings observations, typically sampled at different times and disparate spatial locations, into the model forecast framework to create a dynamically consistent, unified representation of the atmosphere, and eventually of the ocean, to improve subsequent forecasts. To address uncertainty in forecasts, the S-I models must be run in ensembles of 10 to 30 members. The immediate goal is to move from models that resolve the synoptic scales of weather phenomena to models that resolve mesoscale phenomena

at or near the hydrostatic limit such as convective complexes and gravity-wave trains and to move closer to cloud-resolving scales at which different parameterization approaches are required.

C) Forecasting severe weather events: Improving forecasts of events such as hurricanes, tornadoes, and hailstorms has long been a focus of weather research because of the need to reduce the loss of life and property. Because thunderstorms are small-scale, rapidly evolving features, locally at least, models must have fine spatial scales and frequently assimilate data from observations. Data assimilation is an extremely important and computationally expensive component of a complete mesoscale prediction system. In addition, since decision making requires estimates of uncertainty, an ensemble of forecasts must be run to provide not only an optimal estimate but also statistics about its quality. The next major advance in atmospheric prediction will come with the application of storm-resolving mesoscale models (grid spacings of 1 km or less) that are initialized with observations of comparable resolution, triggered in a manner that responds rapidly to the weather itself. To achieve this advance, while producing forecasts 10 times faster than the weather evolves, will require a 100x increase in sustained performance compared to what is available today on 1,000-processor high-end MPP systems. The ability to produce such forecasts for several parts of the country or world simultaneously would introduce an additional multiplier.

D) Space weather: The increasing reliance on satellites for telecommunications and GPS location, as well as the vulnerability of terrestrial telecommunications and power grids to strong electromagnetic storms in the upper atmosphere, has given rise to the field of “space weather.” Since the main driver for space weather is the activity of the sun, this research involves a range of modeling tasks including eruptive events on the surface of the sun, simulating the solar wind and its interaction with the Earth’s magnetosphere, and modeling the interaction between the magnetosphere and the upper atmosphere.

Aspects of space weather modeling that challenge available computer resources include the complexity of the magneto-hydrodynamic (MHD) equations, the need, in some regions, to account for the discrete nature of charged particles, and the very broad range of space and time scales involved. (Some effects propagate at speeds close to the speed of light.) As in terrestrial weather, data assimilation will also be an important component of space weather modeling. To fully realize space weather forecasting requires both software development and increased hardware capabilities. With techniques that are currently being developed, resolution of key processes that control space weather should require a 100x to 1,000x increase in available sustained computing capability. For simple three-dimensional MHD codes, the longer-term goal of developing effective data assimilation techniques for magnetospheric modeling is likely to require 10 to 100 times as much capacity as presently available, with some increase in capability. With the inclusion of uncertainty forecasting, this requirement would probably increase by another order of magnitude.

HEC Requirements: In addition to the increases in sustained throughput noted above, there are architectural requirements of high-end computing for this type of research. These needs are typified by very large bandwidths between memory and processors and between computational nodes (the “interconnect”). The unavoidably global nature of some steps in these calculations also introduces a requirement for a high degree of internodal connectivity (high bisection bandwidth). These needs cannot be met simply by the operation of Moore’s Law on existing

COTS-based parallel systems since Moore's Law does not really apply to such things as memory speeds. Instead, these requirements indicate a need for architectural innovations. High-resolution climate and weather models also generate a very large amount of data (at least a terabyte per modeled day) and hence require very efficient parallel I/O and large local storage. Better scientific analysis of the model output benefits from more capable visualization and data-mining tools. When used in an operational mode, high-resolution weather models must also ingest large amounts of real-time data and return large quantities of output to the user (possibly another model running on a different system), necessitating high-bandwidth network links. In addition, considerable software engineering resources are required to convert very complex climate codes to run on each new architecture in a timely and effective manner.

APPENDIX A-2. NANOSCALE SCIENCE AND TECHNOLOGY

Introduction: The ability to manipulate matter at the atomic level has led to an amazing realization: Clusters of atoms on the nanometer (nm) length scale (order 10^{-9} meters) often have properties dramatically different from those of bulk matter of the same size. This discovery opens a new frontier for scientific breakthroughs and technological innovation based on science and engineering focused on the nm-length scale. Nanoscale science and technology (NST) aims to discover and exploit the unusual nanoscale properties of matter to enable scientific and technological advance. Imagine the capability to create new materials from individual atoms up to macroscopic dimensions, to engineer new electronic devices that derive their functionality from atoms or molecules – the ultimate scale of miniaturization – as well as to probe DNA and the molecules of life, to develop new drug delivery systems that speed the cure directly to diseased tissue, and to do much more through the deliberate and purposeful manipulation of atoms. Design of advanced materials and nanoscale devices starting from atoms and nanoscale “building blocks” and access to a significant part of the nanoscale scientific frontier lie beyond current capabilities; computers much more powerful than today's are required. Computation is required for progress in NST. It enables quantitative comparison of theories with experiments for verification and discovery and can provide crucial information difficult or impossible to obtain from experiment. Access to usable high-end computation in the short term will yield higher returns on the approximately \$1.7-billion investment to date of the interagency National Nanotechnology Initiative (NNI).

The nanoscale world uncomfortably straddles the more mysterious quantum world and the more familiar world of classical mechanics. A high computational intensity obtains from the need for quantum mechanical calculations, the need for high accuracy, the need to include many atoms, the need to model complex nanoscale systems atoms of different species, the need for long simulation time to include physical processes across time scales from 10^{-15} seconds to 10^{-3} seconds and longer, and the confluence of all of these requirements in the same simulation.

Illustrative Problems: NST encompasses many disciplines and involves yet-to-be-understood phenomena. The representative problems below illustrate, but do not exhaust, the intellectual opportunities of NST.

A) Catalysis: The U.S. petroleum, chemical, biochemical, and pharmaceutical industries are the world's largest producers of chemicals, ranging from “wonder drugs” to paints to cosmetics to plastics to new and more efficient energy sources. Catalysts are key ingredients in 90% of chemical manufacturing processes. Several manufacturers already plan to use nanomaterials in

catalytic surfaces. The structure and function of catalysts lie on the nanoscale. Expensive trial and error currently dominates catalyst design. Computational NST can guide the discovery of new catalysts and optimize their performance through direct analysis of the function of candidate catalysts at the molecular level. High accuracy is required because small errors in reaction energy calculations are magnified, leading to large errors in catalyst efficiency and making economic feasibility difficult to assess.

B) Self-Assembly: Manufacturing materials and devices designed on the nanoscale require the deliberate and orchestrated placement of many atoms. A promising fabrication technique involves nanoscale building blocks having specifically tailored chemical and physical properties so that they automatically “snap together” in a desired way. DNA and other biomolecules may be used as building blocks to direct the assembly of carbon nanotubes, semiconductor quantum dots, and inorganic nanostructures. Developing self-assembly and directed self-assembly techniques requires a comprehensive understanding of the properties of the building blocks and the capability to design new building blocks and predict resulting self-assembled structures. Self-assembly simulations span many length and time scales. Three grand-challenge-scale problems have recently been identified: (1) simulation of the self-assembly of templated nanoporous materials, (2) self-assembly of nanostructures directed by DNA and DNA-like synthetic organic macromolecules, and (3) directed self-assembly of quantum dot arrays and other nanoscale building blocks using physical synthesis and assembly methods.

C) Nanoelectronic Devices: Information technology will be a primary beneficiary of advances in nanoscale electronic device technology. Electronic devices involve the control, manipulation, and transport of electronic charge. Nanoscale device technology is not simply further miniaturization. Electronic device operation on the nanoscale differs qualitatively from that of traditional devices. Design and simulation present scientific and computational challenges. The discrete nature of electronic charge can no longer be ignored. Quantum mechanics is required to describe non-equilibrium transport, tunneling, fluctuation effects, discrete energy levels, and the consequences of interactions among electrons. The structure of nanodevices and motion of their constituent atoms become intimately intertwined with electron transport and device function. With current computational resources, a computationally derived comprehensive picture of nanoelectronic and nano-optoelectronic device operation lies out of reach. Important physics is excluded or included inaccurately for ease of computation. Simulations are often effectively limited to a single device; for comparison, a modern PC requires millions of transistors. Experiments are extremely difficult in this regime, and simulations utilizing quantum transport theory are critical for development, design, and engineering. Computational design is already used in the microelectronics industry. The inclusion of quantum effects in the design process requires high-end computing. Device technology on the nanoscale opens new possibilities of utilizing the spin of the electron in addition to its charge for device operation. The controlled manipulation of electron spin also enables a unique kind of massively parallel computing that can dramatically outperform current computers for specific tasks. The design of many “spintronic” devices and any “quantum computer” requires high-end computing.

D) Interfaces with biology: Large biological molecules such as DNA and proteins lie on the nanoscale, and biological cells are natural examples of functional nanostructures. Understanding cellular processes and other biological phenomena at the nanoscale may lead to new nanodevice innovations. Biological molecules that form the basis of cellular function can be made to function as they do in nature or in novel ways (e.g., to assemble semiconductor

quantum dot arrays). High-end computation can help open the frontier at the intellectual interface of the physical sciences with biology. For example, high-end computation enables microscopic modeling of DNA molecules, DNA repair mechanisms, and the interactions between drugs and DNA. Simulations may open new and more efficient methods for DNA sequencing and help to model cellular processes. Combining inorganic nanostructures and nanodevices with biological molecules and systems opens new directions to develop sensors for potentially harmful biological agents, useful probes of cellular function, and prosthetic devices. Simulation and design of these involves the complexity and computational intensity of high-accuracy modeling of large molecules together with simulation of functional complex nanostructures.

E) New phenomena: NST is broad and encompasses many yet-to-be-understood nanoscale phenomena, structures, and processes. Inquiry in this field is ripe for new fundamental scientific discoveries that alter the course of research. Computation will play a vital role in unraveling myriad scientific mysteries on the nanoscale.

HEC Requirements: NST requires a confluence of disciplines involving the expertise of chemists, engineers, physicists, materials scientists, and applied mathematicians. Meaningful simulation and modeling will often require that well-honed tools from different disciplines be made to work together. The performance of these tools spans several orders of magnitude in terms of the amount of memory and time needed to obtain a solution. Some tools will benefit from more powerful computational nodes, others will benefit from increased bandwidth between computational nodes, and still others will benefit from both. Simulations involving large numbers of atoms or electronic states create large amounts of data and high-speed I/O and high-capacity mass storage are required. The diversity of problems and solution schemes precludes the selection of a single computer architecture as the solution, arguing instead for a suite of heterogeneous computing platforms. Algorithmic development will play at least an equal role in advancing computational capabilities in NST: New algorithms are often enablers for previously intractable calculations. Similarly, sometimes even a modest increase in one aspect of hardware capability can be an enabler for beneficial algorithms. One example: Increasing the memory per node in parallel computer architectures has had a remarkable impact on which codes could be parallelized. Visualization is becoming an increasingly important tool for observing the evolution of defects, structures, and related phenomena. Real-time visualization enables simulation “steering” to initiate and study infrequent events.

Many simulations, such as self-assembly, involve a tradeoff among accuracy, the length of the physical time simulated, the number of atoms included, and the allowed complexity of the electronic states. The number of atoms and physical time simulated limit the physical and chemical processes included in the simulation. Often, a large number of atoms simulated for long physical times is needed to capture the essential physical and chemical processes. Accuracy is limited by the level of physical description of the electrons. The lower the level of description, the lower the accuracy and the lower the computational intensity. Some HEC requirements are evident from examining very large low-accuracy simulations and pushing the limits of the highest-accuracy simulations to improve the reliability of simulation on the nanoscale.

With current computational resources, low-accuracy molecular dynamics simulations in which the electronics are replaced by simplified force laws can be performed for up to several billion atoms on time scales of up to a few nanoseconds. This enables the simulation of a single grain in a nanostructured material at a level of accuracy suitable for many mechanical properties.

Predicting the mechanical properties of nanostructured materials or simulating the assembly of many grains involves the interaction of many such grains, leading to simulations that exceed our current capabilities. A 100x to 1,000x increase in computer performance will extend our abilities to calculate mechanical properties of nanostructured materials and simulate aspects of their assembly.

The most accurate quantum chemistry codes can simulate only a small number of atoms and access physical processes on the shortest time scales. Often such high accuracy (~1 kilocalories per mole [kcal/mole]) is required for at least part of the system to accurately model catalysis or the function of enzymes. An increase in computer performance of nearly six orders of magnitude is required to access even the 10-nm scale with this level of accuracy, for a few hundred heavy atoms. At this level of computational power, Car-Parinello methods become practical for million-atom simulations. While less accurate than quantum chemistry, these density-functional-theory-based methods are much more accurate than molecular dynamics and would enable more reliable simulations in a wide area of NST, including self-assembly, design of functional nanosurfaces, and biomolecules.

For some simulations and areas of research, one can identify small spatial regions where high accuracy is needed to capture important physical processes on small length scales. In the rest of the simulation, lower-accuracy methods can be used. Multiscale algorithms carry out this strategy. While some codes exist, developing multiscale algorithms for NST remains an active and challenging area of research as it often involves coupling classical mechanical and quantum mechanical simulations. Multiscale simulation together with high-end computing enables the simulation of complex nanoscale systems and devices, as well as nanostructured materials.

APPENDIX A-3. LIFE SCIENCES APPLICATIONS

Introduction: The biological sciences have experienced tremendous growth over the past five decades. For the first time, there is the potential to understand living organisms as complex dynamical systems and to use large-scale computation to simulate their behavior in a variety of new ways thought impossible just a few years ago. The challenges involved in predictive modeling of biological processes arise from the large size and long time scales of biochemical phenomena, combined with the need for extreme accuracy.

The time scales of biological functions range from very fast femtosecond molecular motions, to multi-second protein folding pathways, to cell cycle and development processes that take place over the order of minutes, hours, and days. Similarly, the dimensions of biological interest range from small organic molecules through multi-protein complexes, cellular processes, tissues, and to the interaction of human populations with the environment. The linking of biological phenomena at all levels of temporal and spatial scales is leading to the transformation of the separate areas of biology to that of “systems biology.” The new science of systems biology is based on the ability to “read” and understand the complexity of biology, beginning with genome sequences and other sources of high-throughput data, including global experimental strategies, coupled with a detailed understanding of the behavior of proteins individually and in complexes.

Modeling multiple levels of biological complexity with fully coupled scales is well beyond even the next generation of supercomputers, but each increment in the computing infrastructure makes it possible to move up the biological complexity ladder and investigate problems that could not be solved previously. New scientific methods and technologies that integrate physical,

chemical, and biological approaches with information science, mathematics, and computational science are being developed. However, these developments can lead to scientific breakthroughs only if accompanied by dramatic changes in access to high-end computing resources.

Illustrative Problems:

A) Cell Signaling in Cancer Initiation: Biological processes typically involve a complex of a large number of individual molecules interacting with assemblies of very large molecules. At the smallest spatial scale, small- to medium-sized messenger molecules can interact with proteins to lead to cell signaling events that communicate information from the outer membrane of a cell to the nucleus. Disruption of this signaling pathway can have dire consequences as illustrated by the finding that mutations of one enzyme involved in this pathway have been found in 30% of human tumors. A thorough understanding of this pathway is of crucial importance for the development of cancer treatments, as well as for an understanding of how cancer is initiated. Because computer simulations can provide atomic-level detail that is difficult or impossible to obtain from experimental studies, computational studies are essential. However, this requires the modeling of an extremely large complex of biomolecules, including lipid bilayer membranes, transmembrane proteins, and a complex of many intercellular proteins, combined with many thousands of water molecules and solvated counter-ions to balance the charge of the protein complexes. This leads to computational analyses involving on the order of more than five million atoms. While such a simulation would provide the first molecular view of the initiation of the signal transduction pathway, it would require tens and possibly hundreds of teraflops per calculation.

B) Understanding the Structure and Function of Biological Systems: A major question for biology is to understand how genes dictate the structure and thus the function of proteins. Computational genomics is the use of information technology and computers to understand where new genes are and how they are controlled within specific DNA sequences. With the recent advances in gene sequencing technology, the acquisition rate of human and microbial whole-genome sequences has increased dramatically. Assembling and interpreting such data will require new and emerging levels of coordination and collaboration in the genome research community to formulate useful computing algorithms, data-management approaches, and visualization systems.

The complete set of proteins expressed by a specific organism (its proteome) is much more complex than the genome of that organism. For example, there are at least 300,000 proteins encoded by only about 30,000 genes. The use of high-throughput technologies involving mass spectrometry, DNA microarrays, and other technologies offer the promise of rapid and comprehensive identification of expressed proteins and protein complexes, and for uncovering associated regulatory networks. These technologies rely heavily on computational approaches and algorithms to deconvolve and archive the raw data, and for matching mass spectrometry tags to protein databases. Specialized visualization systems will be essential for displaying protein interaction networks, mapping data to pathways, and examining computational results from cluster analysis.

Current simulations of cellular networks are limited to subsets of processes in individual cells or to simple cellular interactions, yet simulations will play a critical role in understanding complex biological processes such as the passage of ions through channels across cellular membranes, multispectral analysis of cellular signaling processes, and metabolic networks. Such simulations will involve “mesoscale” models that include three-dimensional continuum

transport and chemical processing of ions and signal molecules. These simulations are computationally complex, but will be important because quantitative descriptions of these processes will be required both to complete our understanding of the functions of the living world and for many promising future applications.

C) Rational Redesign of Enzymes to Degrade Chemical Agents: Understanding structure-function relationships in proteins will enable the study of factors and biochemical pathways that lead to toxin formation, as well as methods for interrupting these pathways, either for detection and prevention or remediation. Specifically, computational approaches can be used for the rational redesign of enzymes to degrade chemical agents. For example, combined experimental and computational efforts can be used to develop a series of highly specific analogs to the enzyme phosphotriesterase (PTE), an enzyme that can be used to degrade neurotoxins. PTE could be redesigned for optimum activity at specific temperatures, optimum stability and activity in non-aqueous and low-humidity environments or in foam, or for improved degradation of neurotoxins. Another enzyme, acetylcholinesterase (AChE), is a key protein in the hydrolysis of acetylcholine. If AChE is inhibited through a phosphorylation reaction with nerve agents such as DFP, sarin, or insecticides such as paraoxon, severe intoxication and death can rapidly result. It is possible to use advanced computation to design better reactivators of AChE that can then be used as more efficient therapeutic antidotes against nerve agents.

D) Neuroinformatics: The human brain involves complex interactions of biological systems ranging from the levels of molecules and genes to the levels of nervous systems and behavior. Mapping the brain is challenging not only because of the complexity of the organ and its processes, but because of the substantial variance in structure and function of the brain among individuals. The explosion of information about the brain at every level of inquiry has resulted in the necessity of relying on advanced technologies of computer and information sciences to integrate data at all scales of resolution and across disciplines into easily accessed sources and derived information products.

Neuroinformatics combines neuroscience, behavioral, and informatics research to develop databases and neuroscience knowledge management systems. The field requires advanced tools and approaches to enable efficient data sharing and data integration. Research in informatics includes databases, graphical interfaces, information retrieval, multiscale simulation, data visualization and manipulation, and data integration through the development of integrated analytical tools, synthesis, and tools for electronic collaboration. In order for these advanced information technologies to be put to wide use by the neuroscience community, they should be generalizable, scalable, extensible, and interoperable, and be developed in concert with significant neuroscience research.

E) Other Computational Biology Applications: Computational biology presents many opportunities to advance understanding of complex biological systems and predict their behavior. In the area of human health care, surgeons will benefit from virtual reality interfaces to microsurgical tools that will allow minimally invasive surgery. Mathematical models of pressure in arteries will aid in recommendations for surgery, and models of a patient's anatomy will enable surgeons to increase the precision of their operations. Neurobiologists will benefit by being able to map neurons involved in specific behaviors, thereby leading to realistic models of brain functions. Patients will benefit because computational modeling of the complex interactions among chemicals, genes, proteins, their structures and the reaction pathways will enable the design of pharmaceuticals for specific effects.

In the area of fundamental bioscience, simulations will enable greater understanding of ion motions in transmembrane channels, elucidate mechanisms of DNA repair enzymes associated with the health effects of radiation exposure, and lead to ways to increase the efficiency of hydrogenase enzymes that microorganisms use to produce hydrogen. Additional areas for the use of computational biology are currently under study by the Biomedical Information Science and Technology Initiative (BISTI), and a major symposium was held in the fall of 2003. (<http://www.bisti.nih.gov/2003meeting/>)

HEC Requirements: The enormous complexity of biological systems and the difficulty of using information from small model systems to address complex, collective phenomena at large scales require significant advances in theories, algorithms, software, and hardware. Currently available computing resources allow computer simulations of biomolecular systems to be carried out routinely for about 100,000 atoms for tens of nanoseconds. Computer resources will need to increase in power by at least three orders of magnitude to allow microsecond simulations of systems with several million atoms.

The computer simulation of large biomolecular systems on parallel architectures presents a number of significant challenges in terms of memory management, processor speed, interprocessor communication, I/O bandwidth, and storage requirements. First, the size of these systems requires the use of distributed data models. Second, even the least complex description of intermolecular interactions in terms of effective pair potentials requires a significant amount of interprocessor communication of data. There will be an even greater demand for interconnect data traffic for simulations using interaction functions beyond pair-wise additive contributions, such as electronic polarization. Third, each molecular dynamics time step requires several synchronizations of all processors, which leads to significant load-balancing challenges. Fourth, simulation trajectories need to be stored for subsequent analysis calculations and visualization, presenting additional I/O bandwidth, network bandwidth, and data-storage challenges.

Some of the promise of greater computational capability for life sciences is captured in the following list. Estimated application requirements are described in teraflops (Tflops):

- **HEC performance level of 50 Tflops:** Increases the total simulation time of biomolecular systems from the current 1-to-10-nanosecond scale to the microsecond time scale. Such resources will be useful for studying protein-protein complexes, cell signaling, and protein-membrane and mineral-membrane complexes, and for obtaining statistically accurate thermodynamic properties, binding free energies, and reaction pathways.
- **HEC performance level of 250 Tflops:** Allows simulation of large biomolecular complexes in which slow conformational changes determine biological function. Examples include: enzyme catalysis using hybrid classical and quantum mechanical methods, extension toward millisecond simulation times, protein folding and unfolding, membrane transport of simple ions and small molecules, membrane fusion, and vesicle formation.
- **HEC performance level of 1000 Tflops:** Simulation of processes in the millisecond time scale, such as protein folding and membrane transport and membrane fusion processes.

APPENDIX A-4. AEROSPACE VEHICLE DESIGN AND OPTIMIZATION

Introduction: The U.S. economy and national defense are critically dependent on air vehicles. These vehicles cover a wide spectrum in size, speed, and performance including small private craft; medium-size to very large commercial and military transports; high-performance military fighters, bombers, and missiles; and hypersonic aircraft and spacecraft. As the size and speed of these vehicles increase, the computational capabilities needed to provide safe, cost-effective, mission-capable designs escalate exponentially. This is due to both scale effects and the nonlinear nature of the supporting technologies. Effective aerospace vehicle design requires optimal integration of multiple interacting disciplines over widely dispersed time and length scales. These disciplines include aerodynamics, aerothermodynamics, structural mechanics, materials, and acoustics. Since these vehicles often fly over populated land areas, environmental issues must be addressed as well as the safety and optimal use of the airspace.

Simulation models are the primary tools used for aerospace vehicle design. The numerical solution of such difficult engineering and physics problems requires computers that have large bandwidth between processors and memory. In fact, memory and interconnect bandwidth are typically the limiting factors on the scale of simulations.

Illustrative Problems:

A) Sonic Boom Prediction/Alleviation: Reduction of sonic booms to acceptable levels is the key technology barrier for supersonic overland flight. The only feasible approach to overcome this barrier is to tailor the vehicle geometry to produce a specially shaped boom signature that is quieter than the conventional N-wave signature. New work in this area targets the development of a validated prediction methodology for shaped-boom aircraft configurations that lead to sonic boom overpressures of well below 1.0 pounds per square foot (psf). Successful completion of these simulations will provide an essential database that, together with measured data from flight experiments sponsored by DARPA and NASA, will enable the necessary calibration of sonic boom design methodology for supersonic aircraft. Both Reynolds-Averaged and DNS Navier-Stokes codes are required for these simulations.

B) Aircraft Noise: Aircraft noise comes from both the engines and the airframe itself. Major advances in noise and combustion simulations are being achieved through the use of full geometry, detailed physics, and time-accurate three-dimensional simulations. The methods and algorithms used in this area of research are being constantly improved, but are computationally intensive.

C) Aviation Safety and Airspace Utilization: Computational fluid dynamics (CFD) is used very effectively in support of turbulence characterization and quantification of hazard as it relates to commercial aviation. The current numerical simulation is carried out with the Terminal Area Simulation System (TASS). The TASS model consists of equations for momentum, pressure, temperature, and water substance (e.g., clouds, rain, and snow). Since the modeling of atmospheric sensors is so accurate, these simulations can be used to generate data sets to be used as “certification tools” for Federal Aviation Administration (FAA) certification of airborne sensors for turbulence hazard prediction.

D) National Airspace Utilization: A leading barrier to improved airport efficiency is the decay time of aircraft wake vortices. CFD simulation can be used to model these effects. Such tools can be used in the development of better prediction models for determining aircraft spacing.

E) Computational Stability and Control: The equations that describe the motion of an aerospace vehicle are nonlinear and coupled. This coupling leads to induced forces and moments – such as side force due to yawing rate. These forces and moments are normally referred to as stability derivatives (SD). Historically, SDs have been obtained from wind tunnel and flight test data. However, it is now reasonable to begin determining these terms computationally. In addition, CFD has the capability to calculate SDs under conditions that are not possible to test for reasons of safety or limitations of facilities. Benefits are substantial. Computational methods have been successfully exploited for aircraft performance predictions, saving by one estimate \$1 billion in development costs. Similar or even larger savings are anticipated as computational stability and control methods mature. However, at least 50 individual solutions would have to be run to obtain the necessary database for SD predictions.

F) Computational Chemistry for Hypersonic Propulsion: The design of air-breathing hypersonic engines (SCRAM jets) requires modeling high-speed chemically reacting flows. This is accomplished using CFD simulations of high-speed reacting flows including both the fluid mechanics and combustion chemistry.

G) Unsteady Separated Flows: Unsteady separated flows occur in many locations on an airframe. One of the more important areas is the high-lift system (such as leading edge slats and flaps). Recent progress with an unsteady DES code has led to significant results on a high-lift system with and without flow control devices.

H) Supersonic Laminar Flow Control: Laminar flow control (LFC) is a key technology for long-range supersonic flight. Cascading benefits for LFC include drag reduction, reduced weight, increased range, lower sonic boom, and reduced noise levels and pollutant emissions.

I) Reentry/Descent of Hypersonic Vehicles: Hypersonic vehicles that reenter the atmosphere from space must release the energy they gained in attaining orbit. Most of this energy is dissipated as heat. Thus the computational simulation of such events requires hypersonic aerothermodynamic (ATD) tools as well as conventional CFD at lower Mach numbers. These ATD tools not only capture the intense heating environment but must also take into account multiple gas species and chemically reacting flows. Normal practice for these efforts is to compute three total trajectories for each candidate configuration – one nominal and two off-design cases. This amount of computation would require about five weeks on a 10-teraflop computer. If launch vehicle simulation is also considered, simulation of a vehicle in ascent will require 100-teraflop-level computations. This level of physical modeling may enable vehicles that are much lighter than today's, leading to lowered cost to orbit.

J) Integrated Multidisciplinary Optimized Design: Robust design optimization is an interdisciplinary area of research that seeks to enable rapid and early integration of high-fidelity nonlinear simulations and experimental results in multidisciplinary analysis (MDA) and design optimization (MDO), thereby reducing cost and development time in aerospace vehicle design. MDA/MDO involves the analysis and robust design optimization of aerospace systems governed by interacting simulations, described by systems of nonlinear, time-dependent, coupled partial differential equations. Current MDO methods employ a simple indirect coupling

using what are called response surfaces. Systematic design optimization methods with high-fidelity analyses and mathematically rigorous adaptation strategies will lead us to radically new designs, difficult or impossible to find with conventional experiential approaches. These developments will reduce design cycle time, risk, and cost; increase robustness of designs; improve system performance with ensured reliability; and enable the designer to assess designs in other than nominal conditions.

HEC Requirements: The use of HEC-based simulation in aerospace vehicle design is characterized by the need to couple a large number of multidisciplinary simulations as part of the design process with turn-around times that are of the order of hours. To explore a sufficient portion of the design space, one typically needs to run on the order of 50 simulations for each part, subsystem, or system. A realistic model may contain 20 million points, making a single analysis with sensitivity analysis computationally intractable with current computational systems. For design to become computationally feasible, the turn-around time per single analysis on a grid of 20 million points will have to be reduced to one hour or less. Systems that could respond to such requirements would necessarily have hundreds of terabytes of memory and at least a petaflop of sustainable computational power. The key challenge to sustaining the computational output of the processors is the memory and interprocessor bandwidth and latency of the system.

APPENDIX B. DISCUSSION ON TIME TO SOLUTION

The real purpose of having supercomputers is to solve problems. Just as purchase cost does not equal total cost of ownership, execution time does not equal time to solution. Time to solution encompasses the time required to code the algorithms, load data and offload results, analyze output, and validate the software, as well as the execution time itself. In many if not most cases, the execution time is not the largest component of time to solution.

There is widespread agreement that trends in both hardware architecture and programming environments and languages have made life more difficult for scientific programmers and contributed to increased time to solution. Early supercomputers featured a single processor with relatively balanced computation speed and memory access times. With all memory equidistant from the processor and no need to synchronize among multiple computation engines, the execution model viewed by the programmer was quite straightforward. More recently, to accelerate beyond the limitations of only a Moore's law advancement in the speed of the microprocessor, we build our largest machines out of hundreds to thousands of processors. Programmers are now faced with a multilevel hierarchy of memory latencies and bandwidths, and the need to orchestrate the synchronized execution of large numbers of parallel execution threads.

While the hardware architecture of computers has changed dramatically since the 1960s, the programming languages have hardly changed at all. Fortran and C, developed for 1960s mainframes and 1970s microprocessors, respectively, remain the dominant languages. Hence the level of the discourse by which humans communicate with machines has not adapted to the increasing complexity of the execution model. Furthermore, while programmers of commercial applications can often take advantage of the rapidly improving clock speeds by abstracting away many of the details of the underlying hardware, programmers of the most challenging scientific applications must know the hardware details intimately in order to extract a sufficient percentage of the machine's potential performance to render their problem tractable in a reasonable time.

The rapidly evolving nature of computing platforms presents further difficulties for the developers and/or users of simulation codes. A large number of very useful scientific codes were not developed for present-day parallel architectures. Some means must be found to accommodate these codes on existing parallel systems and to ease their transition to future architectures. With multiple architectures available as targets, and to be generally applicable both now and in the future, new simulation codes must be built on virtual system abstractions and with lower-level software providing the hooks to the specific system to be used. Some efforts have been made in which programmers have started from scratch in order to develop codes, which are designed for parallel systems using modern principles of software design: NWChem, NAMD, and WRF are some examples. These codes do produce performance increases on parallel systems compared to their older equivalents, but for some types of applications, the efficiency remains sensitive to the architecture. Significant development costs (people and time) are involved, affecting the overall time to solution.

This is not to say that no progress has been made for parallel system software. The development of PVM, followed by MPI, has provided the programmer some standardization of interface for parallel codes, at the cost of significant per-message overhead. The subsequent development of OpenMP for use within a cluster coupled with MPI across clusters extended that “standard.” But it was at the cost of presenting yet another layer of complexity to the programming model, which the programmer must take into account if performance is the goal. Some progress has been made for libraries useful to simulation codes. Building on the work of the BLAS, LAPACK, and ScaLAPACK, PETSc has packaged most of the solvers for nonlinear systems into the generally accepted standard.

However, much remains to be done in order to make both current and future HEC systems more efficient. Compilers map the underlying hardware architecture into a programming model that can be mastered by humans. To unburden the programmer, significant advances in compiler technology are needed; low-level hand coding is not only time consuming, it is neither general nor portable. Debuggers are required for development of new codes for parallel systems and for migration of legacy codes. Links to compiler information are essential for debuggers to generate traces. Better ways to express parallelism at the highest level, via languages, are essential for ease of use and efficiency, both for the developer and for the code itself. Finally, even when codes are run, “time to solution” includes collecting, storing, digesting, and understanding outputs. One of the neglected areas of modern high-end computing is the ability to stream computed data to storage systems, i.e., parallel I/O. Finally, advanced visualization tools are needed to provide insight into results from complex multidimensional, time-dependent simulations in complex geometries.

All of these realities argue strongly for a program balanced between hardware and software. The last ten years of high-end computing have demonstrated a fundamental difference between hardware and software. While Moore’s law has provided dramatic increases in processor speed, developments in systems software and programmability have come much more slowly. Producing software is still a people-intensive activity. There is no Moore’s law for software.

APPENDIX C. HEC R&D TECHNOLOGY AREAS

Hardware

Microarchitecture – The most recent version of the International Technology Roadmap for Semiconductors³⁶ predicts that microprocessors with 1.5 billion transistors will be introduced by 2010. (This is a five-fold increase.) However, common business economics will continue to drive microarchitecture design to optimize performance for high-volume commercial applications. Optimizing chip functionality to support HEC performance requires different chip architectures and subsystems that may not have immediate impact in the commercial market. The challenge lies in utilizing rich transistor budgets of the future in innovative designs to more effectively address the requirements of high-end applications. Opportunities for microarchitecture research include dynamic optimization of hardware, latency-tolerant mechanisms, introspection, fault isolation and recovery, as well as several others.³⁷ In addition, non-traditional and reconfigurable processors (e.g., architectures based on FPGAs, PIMs, or ASICs incorporating licensed IP cores) promise substantial improvements in performance, cost, and thermal generation for many HEC applications.

Memory technologies – The clock rates of commodity microprocessors have risen steadily for the past fifteen years, while memory latency and bandwidth have improved much more slowly. For example, Intel processor clock rates³⁸ have doubled roughly every two years, from 25 MHz in 1989 to over 3 GHz today as a result of Moore’s Law. Memory latency has improved at only 7% per year. Thus a rule of thumb for HEC systems now becomes “Flops are free, but bytes are expensive.” As a result, HEC application performance often falls well below 10% of peak³⁹, partly because processors stall while waiting for data to operate on. Memory latency is not perceived as a problem for consumer applications, which typically spend most of their time waiting for human response. It is also not a large problem for large commercial applications, which can easily divide the many simultaneous processes across additional nodes requiring little communication. Attacking the processor-memory bottleneck for large-data, high-communication, or random access HEC applications will require research and development. This area of research is expected to overlap with investigations into microarchitectures and non-traditional processors because of continued increases in the scale of integration afforded by the advances in silicon processing technology. Promising areas of R&D include design of memory subsystems that support latency-tolerant processors, such as morphable processors, tiled architectures, processors with larger caches, faster memories, fast memory I/O drivers and buses, intelligent or adaptable caches, and smart memories.⁴⁰

Interconnect technologies – The processor-memory bottleneck is only part of the communication challenge in HEC systems. Communication latency and bandwidth limits must be greatly improved across the entire hierarchy of interconnects: intra-chip, chip-to-chip, board-

³⁶ <http://public.itrs.net>

³⁷ Dally, William, *The Last Classical Computer: ISAT Study*, Stanford University, August 2001.

³⁸ <http://www.intel.com/labs/eml/patinterview.htm>

³⁹ This is the theoretical maximum number of operations per sec and is never actually achieved.

⁴⁰ http://velox.stanford.edu/smart_memories

to-board, cabinet-to-cabinet, and system to peripheral (e.g., storage and network). R&D is needed to develop new interconnect technologies (such as optical-based, high-end Infiniband) and architectures (fabrics) that scale to the largest HEC systems, with much lower latency and much higher bandwidth than are commercially available.

Packaging and thermal management – Modern processors and other integrated circuits (ICs) are driving advancements in commercial packaging and thermal management technologies. HEC systems drive these technologies even further, with requirements for the most aggressive integration (to minimize communication latency) and thermal dissipation (to remove heat generated by more tightly packed and higher-frequency processors and ICs). The July 2002 *Report on High Performance Computing for the National Security Community* lists a number of IC packaging and thermal management technologies that require R&D to address the demands of high-end computing. Packaging technologies include multichip modules, stacked ICs, wafer-scale integration, new board and module technologies, new assembly technologies, and others. Thermal management of future HEC systems will require R&D for new liquid, air, and cryogenic cooling technologies beyond what the market will produce on its own.

I/O and storage – Scalable high-performance I/O mechanisms, file systems, and storage devices are essential to high-end computing. Terascale and petascale applications generate massive amounts of data that must be stored, managed, and assimilated. Storage technology trends are resulting in significant increases in storage capacity, but bandwidth and latency improvements occur much more slowly. Major improvements in scalability throughout I/O and storage are required to keep pace with data generation capability and storage capacity. This is an area that has a long history of relatively weak government support at the high end, and significant investment is required to foster timely improvement. Otherwise, the risk of write-once, read-never data is substantial.

Software

Operating systems – A key roadblock to the realization of HEC systems with high sustained performance is the extreme difficulty of scaling commercial operating systems to effectively utilize large numbers of processors (greater than 64) as a single system. Most commercial operating systems cannot maintain cache coherency and a single-system image (SSI) even at this modest scale, and future HEC systems will contain 100,000 or more processors. R&D is needed to scale operating systems to support much larger SSI, which is beyond the interest of the commercial computer industry. This would allow application developers to use relatively painless parallel programming models (e.g., OpenMP and MLP) as compared to MPI, and it would push system vendors to develop very fast and efficient interconnect fabrics, which would help address the memory bottleneck as well. Operating systems should also address reliability, availability, and serviceability (RAS), automatically handling software and hardware faults. In general, run-time capabilities also need substantial enhancement to support high-end computing, not just to address RAS, but also for automated data movement (e.g., for DSM), load balancing, and scheduling. Software libraries, languages, and compilers all need special attention in the high-end computing context.

Compilers – HEC vendors supply a base set of compilers (e.g., Fortran, C, C++) with a HEC system. However, these compilers are developed and optimized for commercial systems and applications, and are typically not well optimized for the HEC systems of interest to the scientific and engineering communities. Furthermore, there is usually no support for languages

that support the DSM programming model, such as UPC or Co-array Fortran. The July 2002 *Report on High Performance Computing for the National Security Community* nicely described the difficulty of building wide vendor support and user adoption of languages that are preferable for high-end computing. Thus, R&D beyond what vendors will undertake is needed for HEC compilers to maximize sustained and parallel application performance to the full size of the system. These compilers must be fully aware of the HEC system processor, memory, and interconnect architectures. In the long term, high-level languages, compilers, and libraries for performance portability across an array of HEC systems, including reconfigurable and non-traditional processor-based systems, are also HEC challenges with little commercial value, but with a potentially major benefit to HEC time to solution.

Programmer tools and environments – The current class of HEC systems is extremely difficult to program for or port applications to, such that without a heroic effort, applications rarely achieve more than a few percent of the peak capability of the system. The development of better programming tools (porting, debugging, scaling, optimization) would substantially improve this situation, enabling true performance portability to new architectures. These tools must become vastly easier to use, totally seamless (integrated into an IDE), completely cross-platform, and highly efficient when applied to applications running at the full scale of the system. Industry representatives have stated that the development of such a totally seamless environment is well outside their scope and will take additional R&D by the high-end community. Promising technologies include tools that are visual, language-free, highly automated and intelligent, and user-friendly, and that address more than just CPU performance. Such high-level programming tools should eventually include support for non-traditional HEC systems, for example, based on reconfigurable FPGA processors or PIMs.

Algorithms – Algorithmic improvements can yield performance improvements at least as significant as those realized by hardware and architectural advances. For the past decade or more, algorithmic research has focused intensely on the capabilities of a single class of architecture – clustered SMP systems. Emerging architectures will be radically different either in scale of implementation or in fundamental computational characteristics that must be taken into account in the development of next-generation algorithms. It is imperative that research in algorithms that can better exploit the current architectural solutions as well as future approaches be conducted in parallel with hardware, software, and systems research.

Systems

System architectures – Recent HEC investments in the U.S. have focused on cluster-based systems almost exclusively. These cluster-based architectures are not well suited for all applications. Indeed, there are national priority applications that might better be mapped onto systems having very different architectural features (for example, significantly improved communication to computation ratios). The current HEC focus on clustering hundreds of small nodes, each with a separate OS, results in poor parallel efficiency, generally below 10% and sometimes lower than 1% of peak on some applications. There is a need for processors matched to interconnect fabrics that perform well on global data access. Larger shared-memory nodes can dramatically reduce programming effort and increase parallel efficiency. The ability to deliver sustained performance on HEC systems in 2010 will depend to a large degree on the successful achievement of data choreography. This is another reflection of the reality that data movement is becoming more of a limiting factor than arithmetic operations for sustained performance in HEC systems. Successful system architectures will feature high-bandwidth

interconnects with minimal latency, tightly integrated network interfaces, large numbers of outstanding memory references, and control over data placement and routing. Effective node architectures (e.g., larger shared-memory/single-system-image nodes) are needed to facilitate scalability and ease of use.

System modeling and performance analysis – There is a need for system/architecture modeling and simulation to allow users to evaluate potential configurations against application requirements and to determine properly “balanced systems” for their applications. The systems being developed today and those required to satisfy petaflop needs will be so complicated that simulation and modeling of these systems, at various levels of detail, will be required to explore technological alternatives. The R&D to support this desired application performance has not yet been undertaken. These models could also be used to establish a system performance signature that could subsequently be used in assessing a system’s viability in the context of application requirements. Finally, these system models could be used to explore parallel programming paradigms and models of I/O performance.

Reliability, availability, and serviceability (RAS) and security – Substantial increases in size and complexity are expected as systems grow to the petascale level and beyond. It is entirely likely that the rate of parts failure will become high enough relative to application execution time that the occurrence of faults during execution will be the norm rather than the exception. This will most likely require an entirely new approach to system reliability and fault management. Focused high-end computing research in RAS, fault tolerance, and system self-awareness is necessary. This research should be closely coupled with system simulation and modeling to facilitate the development of new approaches and yield deep understanding of petascale RAS issues. Scalable security mechanisms are essential to maintaining the integrity of HEC systems, and the complexity of new architectures will have to be accounted for in focused security research.

Programming models – Alternative programming paradigms to MPI (or MPI-like) are needed, including methods for evaluating their efficacy, as are methods to accelerate collective operations, perform remote DMA, and offload work to a NIC. Advances in new programming models must complement advances in new architectural approaches. Systems consisting of hundreds of thousands of processors must be abstracted if any productivity improvements are to be realized. This effort also requires that sufficient attention be paid to properly socializing the concepts and ideas to the entire high-end computing community.

APPENDIX D. DISCUSSION ON SYSTEM SPECIFICATIONS AND APPLICATION REQUIREMENTS

While the needs vary substantially from application to application, it is clear that all science and engineering domains would benefit from significantly larger, scalable computer architectures that exhibit an improved balance in processor performance, memory bandwidth/latency, communications bandwidth, and improved programming environments. However, simply increasing the size of massively parallel architectures based on current COTS technologies will not effectively address the capability needs of science and engineering research in the coming decade.

Many applications have a critical requirement for high-end computers with high random access global memory bandwidth and low-latency interprocessor communications. While such applications are well suited to symmetric multiprocessors, they are not readily programmable on commodity clusters. Scalable MPP and cluster systems, while providing massive amounts of memory, are inherently more difficult to program. Even emerging parallel vector architectures, such as the Cray SV2/X1, will also be more difficult to program than previous parallel vector systems because of their non-uniform memory access rate.

Numerous attempts are currently under way to retool codes in application areas such as stockpile stewardship, global climate modeling, computational fluid dynamics, local and regional weather forecasting, aircraft design, cosmology, biomolecular simulation, materials science, and quantum chemistry to run more efficiently on MPP architectures, simply because they are the most plentiful systems currently available for high-end computing. These efforts, while they have resulted in more scalable codes in the short run, have diverted attention away from the development of systems that provide high-bandwidth access to extremely large global memories. While some of these research areas stand to benefit significantly from factors of 10 to 100 increases in MPP scale, in general they will still not be able to capture more than about 3% to 20% of the peak processor performance on these systems.

Simulations in astrophysics, climate modeling, computational biology, high-energy physics, combustion modeling, materials research, and other scientific applications will routinely produce data sets over the next few years that range in size from hundreds of terabytes to tens of petabytes. It is estimated that the high-energy and nuclear physics community will have several large-scale facilities by 2006 that will generate tens of petabytes of data per year. The full promise of high-end computing will not be realized without adequate storage, file systems, and software to manage, move, and analyze these datasets effectively. A well-balanced HEC facility demands a smooth flow of data to appropriately large storage and file systems and the technology to efficiently search and mine the data during subsequent analyses. Close coupling between scientific data management and other technologies, such as metadata cataloging, data mining, knowledge extraction, visualization, and networking, must be improved. Current technologies for transferring data from MPP systems to external disk or tape systems will be inadequate by factors of 5 to 50, and our current ability to move even terabytes of data over a network for analysis is severely limited by existing communications bandwidth. Significant investments must be made in storage infrastructure and storage technology over the next several years to deal with the tsunami of scientific data that will be produced.

APPENDIX E. AN INITIAL DISCUSSION ON GOVERNANCE OF HIGH-END COMPUTING

Carrying out the HEC Revitalization Plan requires a governance structure that can achieve the overarching goals and avoid current weaknesses. Federal management and budget practice requires that HEC activities be executed as a collection of either individual agency or multiagency programs, so the governance structure must complement or supplement these parameters. With respect to governance, an instructive lesson that emerged during the Task Force’s examination of previous HEC planning activities was that good planning is a necessary, but not sufficient, condition for success. Strong leadership, effective management, and coordinated implementation are also necessary for a successful initiative. Therefore, in order to advance senior-level discussions on governance, the following strategies are offered for consideration and evaluation.

GOVERNANCE STRUCTURE

A governance structure considered by Task Force is displayed in Figure 3 on page 66. In this model, a HEC Interagency Program Office (IPO) is responsible for the interagency coordination that is crucial for the success of the HECRTF Plan. Coordination would occur in the areas of long-range planning, budget preparation, HEC program review, R&D integration, HEC access, procurement, and management of the set of Leadership Systems by the agency (or agencies) responsible for them.

The IPO, preferably integrated with the National Coordination Office for Information Technology Research and Development, would have a full-time director on detail from a participating agency. Senior management from each participating agency would formally designate agency members to the Office. Office “membership” would include program managers and technical staff serving on either a full-time or a part-time basis. The Office would report to a committee of the NSTC, either directly or through the NITRD Interagency Working Group, composed of senior agency executives representing both end-user communities and the HEC R&D community. The Office would be assisted by a Science and Engineering Team comprising end-user representatives from participating agencies. The Science and Engineering Team would ensure that the needs of end users are reflected in HEC plans and would relay the opportunities of HEC technology back to the user communities. Within their agencies, Science and Engineering Team members would encourage partnering between user communities and the developers and providers of high-end computing.

HEC Interagency Program Office

Coordination and management practices of the Office would depend on the nature of specific program activities. Agency membership in the Office denotes an official relationship: members represent agency activities and needs in an authoritative manner with the support of senior leadership. Members of the Office would coordinate pre-budget planning activities, meeting approximately monthly. The Office would develop metrics to ensure that ongoing and new-start

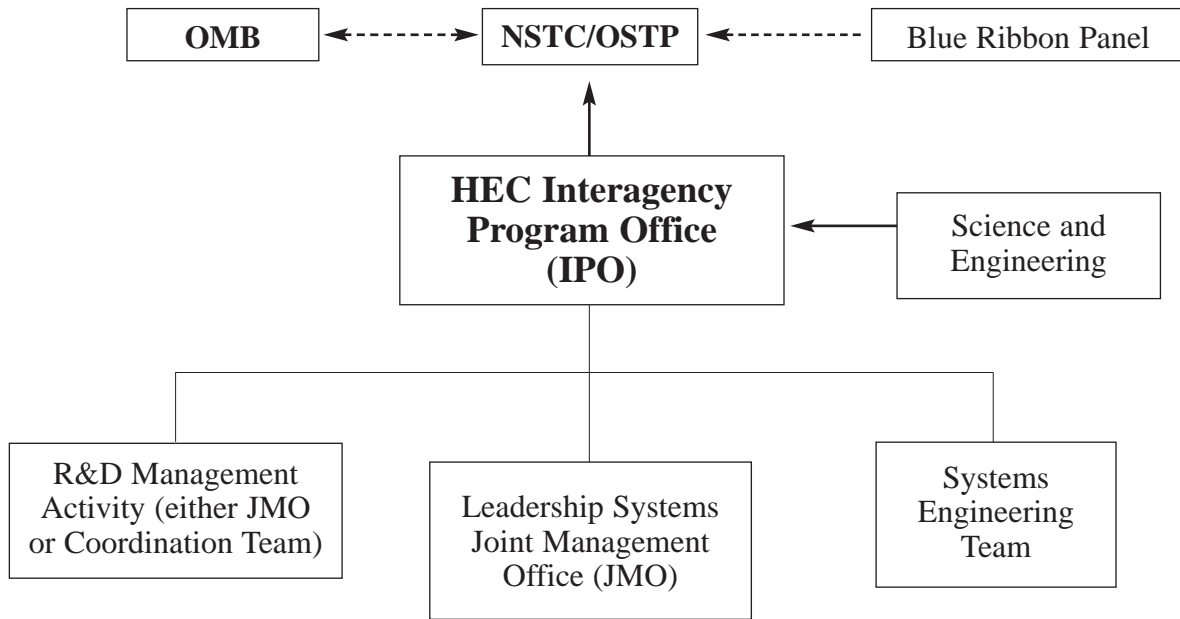


Figure 3: Proposed Federal HEC Governance Structure

programs include appropriate interagency coordination. This should include joint preparation of program solicitations, evaluation of proposals, program reviews, and roadmapping activities.

Coordination and Management of the R&D Portfolio

Management and coordination of the R&D portfolio proposed in this Plan present particular challenges. The Task Force has developed two potential options for implementation. The first is an R&D Coordination Team; the second is an R&D Joint Management Office.

Under the first option, agencies conducting HEC R&D activities would first coordinate plans before budget submissions, and each agency would then submit budget plans for HEC R&D activities through its own processes. Each agency would commit to components of the R&D plan and be held accountable via MOUs with the IPO. Controls provided through oversight by the IPO, OMB, OSTP, and the NSTC would serve to enforce adherence to the pre-budget plan.

This model would permit agencies to jointly manage R&D activities on a voluntary basis. For example, the national security community could manage its HEC R&D activities under a Joint Program Office, as proposed by the DoD “Report on High Performance Computing for the National Security Community.” This model is workable if there is upper-level support for HEC R&D investments and significant cohesion and cooperation among the participating agencies.

Under the second option, the Task Force would investigate establishment of a Joint Management Office (JMO) to oversee the R&D portfolio. The Office would conduct integrated planning, solicitation preparation, project selection and execution, and progress reviews, subject to approval by the participating agencies. Such an option enables activities that are essentially multiagency and within the scope of this Plan, but beyond the resources or expertise of

individual agencies. The advantage of such a model is that project decisions and management authority are highly visible and achieve a critical mass to move the field forward rapidly.

Management of Leadership Systems

Each Leadership System would require a strong integrated team. A JMO or equivalent management approach that includes users and experts from across government in all the critical technical areas, as well as a strong engineering team, would be important to the success of each multiyear effort. These projects, by their very nature, produce first-of-a-kind or even one-of-a-kind systems that may or may not be commercially viable at the scale at which they are developed. The Task Force anticipates, however, that the hardware and software subsystems developed would have a powerful impact on commercially available products. For that reason, the Leadership efforts are likely to require coordinated government and industry activities. New multiagency procurement approaches would be needed to create appropriate management and funding structures. Moreover, the substantial investment required by the Leadership efforts necessitates that these systems be managed as a national resource, analogous to an x-ray or synchrotron light source or neutron source. A working group would be formed with the task of adapting the “Cooperative Stewardship” model⁴¹ for application to Leadership and other HEC systems.

Strategic Projects

Under the HEC Interagency Program Office, ad hoc teams would be formed to attack strategic issues in high-end computing. One such area is systems engineering as applied to high-end computing. Advances in this area would provide insight into and understanding of the trade-offs required to advance R&D and procurement investments. The Task Force’s proposed projects for benchmarking, total cost of ownership, roadmapping, and performance modeling represent elements that the Systems Engineering Team would coordinate and execute.

Roadmaps and Outreach

The semiconductor industry offers a model for the use of roadmaps to change the behavior of an entire industry that may be usefully studied for possible application to high-end computing. During the 1980s, the U.S. semiconductor industry realized that it needed to change to respond to Japanese competition. With matching government funding, the U.S. manufacturers created Sematech and related organizations to improve the process for developing semiconductor technologies. Through Sematech, the manufacturers worked out a more sophisticated strategy for determining when to compete with each other and when to cooperate. The areas they identified for cooperation were pre-competitive, mostly research and process development, and had previously been treated as proprietary. Lessons from that experience, including the role of government-funded research, may have application here.

The IPO would use a similar planning approach to build consensus and sustain momentum in high-end computing. End-user roadmaps would track the needs of applications and synthesize common requirements for HEC technology. Technology roadmaps would assure alignment of technology development with end-user roadmaps; they would guide Federal HEC R&D

⁴¹ National Research Council, *ibid.*

programs and influence industry investments. The Office would encourage broad and inclusive participation in roadmapping activities by academia, industry, government and national labs, and defense contractors, as well as by Federal agencies themselves. Using these roadmaps, the Office would assist participating agencies to develop appropriate R&D programs.

While the IPO would achieve outreach primarily through broad involvement in roadmapping, other activities such as workshops, information meetings, demonstrations, and presentations would be used as vehicles for exchanging information throughout the community.

Program Evaluation

Review of Federal high-end computing activities would be provided by two mechanisms. First, the IPO would conduct internal reviews of participating agency programs and budgets to determine consistency with agreed-upon plans. Review findings would be reported to OSTP, OMB, and the NSTC. Second, an external review would be conducted by a blue-ribbon panel every two years to assess the HEC R&D portfolio and the Federal strategy to address user needs for capability, capacity, and accessibility. External reviews would focus on high-level issues related to effectiveness, balance, progress along technology roadmaps, and interagency teaming. Review findings and recommendations would be reported to the Director of OSTP and the Director of OMB. The Director of OSTP would work with senior-level agency executives to resolve issues identified via these external reviews.

APPENDIX F. LIST OF ACRONYMS

ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
ASCR	Advanced Scientific Computing Research
ASIC	application-specific integrated circuit
ATD	aerothermodynamic
BISTI	Biomedical Information Science and Technology Initiative
BLAS	basic linear algebra subroutine
CAF	Co-Array Fortran
CFD	computational fluid dynamics
CICT	Computing, Information, and Communications Technology
COTS	commercial-off-the-shelf
CPU	central processing unit
DARPA	Defense Advanced Research Projects Agency
DES	detached eddy simulation
DMA	direct memory access
DNA	deoxyribonucleic acid
DNS	direct numerical simulation
DOC	Department of Commerce
DoD	Department of Defense
DOE	Department of Energy
DOE/NNSA	Department of Energy/National Nuclear Security Administration
DOE/SC	Department of Energy/Office of Science
DSM	distributed shared memory
EPA	Environmental Protection Agency
FAA	Federal Aviation Administration
FLOPS	floating-point operations (per second)
FPGA	field-programmable gate array
GF	gigaflops
GPS	Global Positioning System
HEC	high-end computing
HECRTF	High-End Computing Revitalization Task Force
HPC	High Performance Computing
HPCC	High Performance Computing and Communications
HPCMP	High Performance Computing Modernization Program
HPCS	High Productivity Computing Systems
IC	integrated circuit
IDE	integrated development environment
IHEC	Integrated High-End Computing
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
I/O	input/output
IP	intellectual property
IPO	Interagency Program Office

IT	information technology
IT R&D	information technology research and development
IWG	Interagency Working Group
JMO	Joint Management Office
JPO	Joint Program Office
LAPACK	Linear Algebra PACKage
LES	large eddy simulation
LFC	laminar flow control
LINPACK	LINEar algebra PACKage
LTD	live test demonstration
LTOP	lease to purchase
MDA	multidisciplinary design analysis
MDO	multidisciplinary design optimization
MHD	magneto-hydrodynamics
MLP	memory-level parallelism
MOU	memorandum of understanding
MPI	message passing interface
MPP	massively parallel processing
NASA	National Aeronautics and Space Administration
NCO	National Coordination Office for IT R&D
NERSC	National Energy Research Scientific Computing Center
NGA	Next Generation Architecture
NGI	Next Generation Internet
NIC	Network Interface Card
NIH	National Institutes of Health
NIMA	National Imaging and Mapping Agency
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development
NNI	National Nanotechnology Initiative
NOAA	National Oceanic and Atmospheric Administration
NRO	National Reconnaissance Office
NSA	National Security Agency
NSF	National Science Foundation
NST	Nanoscale Science and Technology
NSTC	National Science and Technology Council
ODDR&E	Office of the Deputy Director Research and Engineering
OMB	Office of Management and Budget
OpenMP	Open specification for MultiProcessing
OS	operating system
OSTP	Office of Science and Technology Policy
PACI	Partnerships for Advanced Computational Infrastructure
PC	personal computer
PETSc	Portable, Extensible Toolkit for Scientific computation
PIM	processor-in-memory
PVM	parallel virtual machine
R&D	research and development

RAS	reliability, availability, serviceability
RLV	reusable launch vehicles
ScaLAPACK	Scalable Linear Algebra PACKage
SciDAC	Scientific Discovery through Advanced Computing
SMP	symmetric multiprocessing
SSI	single system image
SSP	sustained system performance
TASS	Terminal Area Simulation System
TCO	total cost of ownership
Tflops	Teraflops
UPC	Unified Parallel C

APPENDIX G. TASK FORCE MEMBERSHIP

HECRTF Co-Chairs

Grosh, John	DoD/ODDR&E
Laub, Alan	DOE/Office of Science
Nelson, David	National Coordination Office
Szykman, Simon	Office of Science and Technology Policy

Core Technologies R&D Task Group

Abdali, S. Kamal (Co-Chair)	NSF	Greene, Thomas	NSF
Muñoz, José (Co-Chair)	DOE/NNSA	Johnson, Frederick	DOE/Office of Science
Bennett, Herbert S.	DOC/NIST	Hart, Leslie B.	DOC/NOAA
Biegel, Bryan	NASA	Lyster, Peter	NIH
Biswas, Rupak	NASA	Schneider, Barry	NSF
Cassatt, James	NIH	van Rosendale, John	DOE/Office of Science
Culhane, Candy	DoD/NSA	Varman, Peter	NSF
Graybill, Robert	DoD/DARPA		

Capability, Capacity, and Accessibility Task Group

Devaney, Judith (Co-Chair)	DOC/NIST	Hoang, Thuc	DOE/NNSA
Page, Thomas (Co-Chair)	DoD/NSA	Lee, Tsengdar	NASA
Cassatt, James	NIH	Lyster, Peter	NIH
Cherniavsky, John	NSF	Mandula, Jeffrey	DOE/Office of Science
Fine, Steven	EPA	Meacham, Stephen	NSF
Frazier, Njema	DOE/NNSA	Merrill, Michael	DoD/NSA
Gross, Brian	NOAA	Polansky, Walter	DOE/Office of Science
Henry, Cray	DoD/HPCMP	Riches, Michael	DOE/Office of Science
Hess, Daryl	NSF	Rohfling, Celeste	NSF
Hilderbrandt, Richard	NSF	Stevens, Walter	DOE/Office of Science
Hirsh, Rich	NSF	Thames, Frank	NASA
Hitchcock, Daniel	DOE/Office of Science	Thigpen, William	NASA

Procurement Task Group

Turnbull, William (Co-Chair)	DOC/NOAA	Kaluzniacki, Roman	DoD/NSA
Webster, Phil (Co-Chair)	NASA	Koelling, Dale	DOE/Office of Science
Bayer, Paul	DOE/Office of Science	Kreisman, Norman	DOE/Office of Science
Hammer, Guy	DoD/Missile Defense Agency	Lee, Sander	DOE/NNSA
Henry, Cray	DoD/HPCMP	Narcho, Graciela	NSF
Hitchcock, Daniel	DOE/Office of Science		

Integration Task Group

Grosh, John (Co-Chair)	DoD/ODDR&E	Meisner, Robert	DOE/NNSA
Laub, Alan (Co-Chair)	DOE/Office of Science	Muñoz, José	DOE/NNSA
Abdali, S. Kamal	NSF	Oliver, C. Edward	DOE/Office of Science
Brooks, Walter	NASA	Page, Thomas	DoD/NSA
Cotter, George	DoD/NSA	Szykman, Simon	OSTP
Devaney, Judith	DOC/NIST	Romine, Charles	OSTP
Freeman, Peter	NSF	Trinkle, David	OMB
Graybill, Robert	DARPA	Turnbull, William	DOC/NOAA
Johnson, Frederick	DOE/Office of Science	Webster, William P.	NASA

National Coordination Office for Information Technology Research and Development

Howe, Sally	King, Frankie	Kirk, Elizabeth	Matzke, Martha
Miller, Grant	Moore, Virginia	Theiss, Diane	

COPYRIGHT

This is a work of the U.S. government and is in the public domain. It may be freely distributed and copied, but it is requested that the National Coordination Office for Information Technology Research and Development (NCO/IT R&D) be acknowledged.

**National Science and Technology Council
Committee on Technology
High-End Computing Revitalization Task Force**

National Coordination Office for
Information Technology Research and Development
Suite II-405, 4201 Wilson Blvd. Arlington, VA 22230
(703) 292-4873 • <http://www.nitrd.gov>

