

June 1999

NUCLEAR WEAPONS

DOE Needs to Improve Oversight of the \$5 Billion Strategic Computing Initiative





United States
General Accounting Office
Washington, D.C. 20548

**Resources, Community, and
Economic Development Division**

B-282727

June 28, 1999

The Honorable Duncan Hunter
Chairman, Subcommittee on Military Procurement
Committee on Armed Services
House of Representatives

Dear Mr. Chairman:

As requested, this report examines the Department of Energy's (DOE) Accelerated Strategic Computing Initiative. Specifically, it discusses the management of the program, including (1) whether the program is meeting its key milestones and whether hardware and software developments are adequate to date, (2) whether the program is within its projected budget, and (3) what key technical risks the program faces.

As arranged with your office, we plan to distribute copies of this report to the appropriate congressional committees; the Honorable Bill Richardson, Secretary of Energy; the Honorable Jacob Lew, Director, Office of Management and Budget; and other interested parties. We will also make copies available to others on request.

If you have any questions regarding this report, please contact me or Edward Zadjura at (202) 512-3841. Key contributors to this assignment were Linda Chu, Daniel Feehan, and Anne McCaffrey.

Sincerely yours,

Susan D. Kladiva,
Associate Director, Energy,
Resources, and Science Issues

Executive Summary

Purpose

Historically, the United States detonated nuclear weapons as the primary method of validating designs and certifying the weapons as safe and reliable. Since September 1992, there has been a moratorium on testing. To ensure the continued safety and reliability of nuclear weapons, the Department of Energy (DOE), which is responsible for designing and building nuclear weapons, developed the 15-year Stockpile Stewardship and Management Program in 1995 as a substitute for actual testing. The stockpile stewardship program employs a variety of means to ensure weapon safety and reliability, including examining weapons, conducting laboratory experiments and tests, and conducting computer modeling and simulation. The computer modeling and simulation part of the program is known as the Accelerated Strategic Computing Initiative. The strategic computing initiative aims to develop advanced computer models that will simulate nuclear explosions in three dimensions with higher resolution than previous models and with a more complete treatment of the underlying basic physics. The initiative is also developing the world's largest and fastest computers, which may ultimately be able to calculate more than 100-trillion mathematical operations per second. The initiative is expected to cost about \$5.2 billion for fiscal years 1996 through 2004.

Concerned about the status of the strategic computing initiative, the Chairman, Subcommittee on Military Procurement, House Committee on Armed Services, requested that GAO review the management of the strategic computing initiative, including (1) whether the program is meeting its key milestones and whether its hardware and software developments are adequate to date; (2) whether the program is within its projected budget; and (3) what key technical risks the program faces.

Background

Since the dawn of the nuclear era in 1945, the testing of nuclear weapons and state-of-the-art computing have been used together to ensure the performance, reliability, and safety of the weapons. Testing was the ultimate judge of whether a weapon worked and met its design requirements and provided data needed for computer models. Computers were used to perform the massive calculations needed to understand the basic physical processes that take place at the heart of a nuclear explosion and to interpret the results of nuclear experiments and tests, thus providing feedback in the process of designing, building, and testing nuclear weapons. The practical result of the Comprehensive Test Ban Treaty of 1996 is that existing nuclear weapons will be kept longer than planned because new weapon designs cannot be tested and certified as safe and reliable. Faced with these testing restrictions, DOE developed a

new approach to certifying the safety and reliability of weapons in the U.S. stockpile. The computer models and hardware, developed as part of the strategic computing initiative, will be used to identify potential stockpile problems by predicting the effects of aging and the need to replace components or even to retire weapons systems if they become unsafe or unreliable. The existing stockpile of weapons is aging, and many of the designers of those weapons have retired or are approaching retirement. For these reasons, DOE wants to have the computers and models available by 2004 so that the existing cadre of experienced weapons designers will be available to help verify the results of the models.

Results in Brief

Weak management and information processes hamper oversight of the strategic computing initiative. Although initiative managers report that many milestones have been met, the lack of comprehensive planning and progress tracking systems make assessment of the initiative's progress difficult and subjective. Currently, the initiative's strategic plan is out of date, annual plans have been prepared only sporadically, and milestones are not well defined. Furthermore, little information exists to track the initiative's progress or to compare its accomplishments with its milestones. Consequently, it is difficult to determine which of the hundreds of milestones have been met, which are behind schedule, or even which are still relevant, given changes in the initiative.

Program cost estimates have increased substantially. In 1995, DOE estimated that costs for the first 5 years of the initiative (fiscal year 1996 through fiscal year 2001) would be \$1.7 billion. By 1999, estimated costs for that same 5-year period increased to \$2.9 billion. DOE currently estimates that the program will cost about \$5.2 billion for fiscal years 1996 through 2004.¹ Some of the cost increases result from the shift to computer-based simulations, while some reflect weaknesses in DOE's cost estimation.

Developing a computer simulation, or "virtual test" capability, that, in the absence of nuclear testing, can be used to determine whether a weapon system will perform as intended requires overcoming significant technical challenges. These challenges range from developing state-of-the-art hardware and software technologies, to integrating scientific data from weapons physics experiments, to recruiting and retaining staff with the needed technical expertise.

¹The strategic computing initiative is a 15-year program, but because of the 5-year budget cycle, no cost estimates are available beyond fiscal year 2004.

This report makes recommendations directed at improving the oversight and management of the strategic computing program.

Principal Findings

A Comprehensive Planning and Tracking System Is Needed to Assess Program Progress

The strategic computing initiative's planning efforts have been inconsistent and incomplete. Strategic planning documents have not been updated, and annual implementation plans were prepared inconsistently or, in some cases, not at all. The long-term milestones presented in various plans are inconsistent, and no information exists to link annual activities to these milestones. Strategic plans also do not identify the multiple research strategies currently employed to meet many long-term milestones, nor do the plans include key decision points for managing these strategies. Performance criteria for most milestones have also not been defined. The efforts of DOE and laboratory managers to track the progress of the strategic computing initiative have been limited primarily to reporting annual accomplishments, without any systematic tracking of progress towards long-term milestones. As a result, it is not possible to determine whether annual milestones were achieved or to what extent annual efforts contribute to long-term milestones. In response to GAO's request for tracking information, program officials have decided to track and report the program's progress more systematically.

The lack of a system for tracking progress, combined with the lack of defined performance criteria, make it difficult to assess whether the strategic computing initiative is proceeding on schedule and delivering the performance expected. However, it is possible to gain some limited insights through discussions with laboratory officials on individual projects or areas. For example, in the area of hardware development, most contract milestones to date relating to the delivery and installation of computers and related hardware have been met, although not all acceptance tests have been passed.

DOE is not managing the strategic computing initiative as a strategic system. To be designated as a strategic system, under DOE criteria, a project must cost over \$400 million, be an urgent national priority, be high-risk, have international implications, or be vital to national security. The purpose of designating strategic systems is to ensure informed, objective, and well-documented decisions for key events, such as changes

to baseline cost or schedule and to ensure oversight at the highest departmental level. The strategic computing initiative meets all these criteria, has experienced delays in some areas, has had its projected costs increase, and depends, in some cases, on as-yet unknown technologies for success. These characteristics, coupled with demonstrated weaknesses in program management and oversight, make the strategic computing initiative a clear candidate for being designated as a strategic system. According to DOE, it has not designated the initiative as a strategic system because the program is already subject to high-level departmental oversight. However, as discussed above, GAO found serious weaknesses in the program's management and information processes that make it difficult to determine if the program is performing as expected.

The Management and Tracking of Costs Need to Be Improved

DOE's cost estimates for the strategic computing initiative have increased substantially since 1995, when early budget projections were made. Costs for fiscal years 1996 through 2001 have increased from an original estimate of \$1.7 billion to the current \$2.9 billion. DOE's fiscal year 2000 budget request for the strategic computing initiative, which totals \$692 million, is more than double the original fiscal year 2000 estimate made in 1995. Some of the cost increases result from the shift from test-based experiments to computer-based simulations, while some increases are the result of weaknesses in DOE's cost estimation.

Although DOE monitors month-by-month spending at the laboratories, it does not track costs for specific projects. As a result, DOE cannot determine which projects, if any, may be costing more or less than originally planned. GAO has previously noted DOE's difficulty in managing costs and schedules in large projects.

Technical Challenges Are Present in All Aspects of the Strategic Computing Initiative

The development of hardware and software technologies and the necessary infrastructure to support these technologies are critical to achieving the simulation and modeling goals of the strategic computing initiative. The program faces significant technical challenges in all of these areas. For example, increasingly large and complex computers using thousands of processors must be developed and made to operate as a single integrated system at speeds far beyond any achieved to date. The effort to develop software for simulation models on the scale needed to model nuclear weapons requires incorporating massive amounts of data, utilizing increasingly sophisticated problem-solving techniques, and using increasingly larger and faster computers. The President's Information

Technology Advisory Committee recently described software of this scale as being “among the most complex of human-engineered structures.” Furthermore, developing the needed infrastructure, including data storage and visualization technology, will require significant technological improvements.

The successful integration of data from laboratory experiments conducted outside the strategic computing initiative into software models being developed as part of the initiative has been noted by the DOE-chartered Blue Ribbon Panel as another important technical challenge. Data from these experiments and past nuclear tests are critical for demonstrating that the results of the software simulations are accurate. According to program officials, a recent reorganization of DOE offices and the creation of a formal software validation program aim to address this challenge.

Finally, recruiting and retaining qualified personnel is a continuing area of risk, according to strategic computing initiative officials and outside program reviews such as the Chiles Commission.² DOE and laboratory officials have efforts ongoing in many areas to improve the recruitment of staff with the required expertise. Nonetheless, as noted by the Chiles Commission, there is no certainty that DOE’s efforts will succeed.

Recommendations

DOE has chosen not to designate the strategic computing program as a strategic system. Given the strategic computing program’s estimated cost of over \$5 billion; the lack of a comprehensive planning, tracking, and reporting system; and the importance of the program to maintaining the stockpile of nuclear weapons; it is important that DOE improve its oversight and management of this program. Therefore, we recommend that the Secretary of Energy require the establishment of a comprehensive planning, progress tracking, and reporting system for the program and designate the program as a strategic system warranting oversight at the highest departmental level.

Given the substantial increases in the cost estimates for the strategic computing initiative to date, DOE’s weaknesses in estimating costs for the unprecedented scale of development efforts, and the lack of a cost-tracking process for the projects under the initiative, we also recommend that the Secretary of Energy require the strategic computing initiative to adopt systematic cost tracking procedures that will allow DOE managers to determine if specific projects are within budget.

²Report of the Commission on Maintaining United States Nuclear Weapons Expertise, Mar. 1, 1999.

Agency Comments and Our Evaluation

GAO provided DOE with a draft of this report for its review and comment. DOE concurred with part but not all of the report's recommendations. Specifically, DOE concurred with the recommendation on the need to improve its oversight and management of this program and cited changes that it has made or is in the process of making. DOE did not concur with the recommendation to designate the program as a strategic system or on the need to adopt systematic cost tracking procedures.

In agreeing with the recommendation to improve the oversight and management of the program, DOE cited several changes it was making. Specifically, DOE stated that it would soon issue an updated Program Plan that will include detailed specifications for all of the critical program milestones. In addition, FY 2000 Implementation Plans will be issued by September 30, 1999, that will include descriptions of all program elements and complete lists of all milestones. The Department also cited the creation of a quarterly progress tracking mechanism to track program milestones. However, in addition to tracking the program's progress against established calendar milestones, it is also necessary to establish specific technical criteria for what constitutes the successful completion of those milestones. Until DOE completes and publishes its revised Program Plan and FY 2000 Implementation Plans, GAO cannot determine whether the Department has fully complied with this recommendation.

DOE disagreed with the recommendation to designate this program as a strategic system, stating that to do so would duplicate the existing planning, progress tracking, and reporting system. GAO agrees that creating a duplicate tracking system that mirrors the requirements set out by DOE for strategic systems would not be worthwhile. However, as discussed in detail in this report, DOE has not shown that it has an adequate planning, progress tracking, and reporting system in place for the strategic computing initiative. While DOE is making some positive improvements in these areas, the changes are not yet fully in place, and their adequacy cannot be judged at this time. Furthermore, if the changes that DOE is making are adequate to meet the requirements for tracking and monitoring of a strategic system, then GAO cannot understand DOE's reluctance to designate this large and costly program as a strategic system. DOE stated that it has a review process that meets the intent of the Clinger-Cohen Act of 1996. However, GAO reported in July 1998 that the Department's process effectively excluded scientific computers like those being acquired through this program from DOE's normal review channels and places them

within the program offices³. GAO stated that all computers should be included as part of the normal DOE Clinger-Cohen review process.

DOE also did not agree with the recommendation to adopt systematic cost-tracking procedures for the strategic computing initiative, noting that costs are tracked by budget and reporting codes in the Department's Financial Information System. DOE stated that these systems are extended down to individual projects with other funding and cost-monitoring tools that gather more detailed information. As an example, DOE cited a March 1999 analysis of selected projects that identified the commitments and cost status for specific procurements at the project level. GAO does not agree that DOE has an adequate level of tracking at the project level or that the changes it is making will rectify this problem. DOE's current system tracks costs only at the aggregate level and does not allow DOE managers to determine which projects at the laboratories are under or over budget. Furthermore, the "other funding and cost monitoring tools" that DOE uses do not allow the systematic tracking of project costs. DOE also stated that some budgeting flexibility is necessary to capitalize on changes within the high-computing industry. While some budgeting flexibility is necessary in a project of this size and complexity, GAO does not believe that this flexibility should preclude effective oversight of a multiyear program costing over \$5 billion.

DOE's written comments are included in appendix II, and GAO's responses are discussed in chapters 2 and 3 and in appendix II.

³Information Technology: Department of Energy Does Not Effectively Manage Its Supercomputers (GAO/RCED-98-208, Jul. 17, 1998).

Contents

Executive Summary		2
Chapter 1		12
Introduction	The Evolving Role of Computing in Nuclear Weapons	12
	Why Such Large Computers Are Needed	14
	ASCI Program Overview	15
	Objectives, Scope, and Methodology	17
Chapter 2		19
A Comprehensive	Comprehensive Planning Is Needed	19
Planning and Tracking	Systematic Progress Tracking and Reporting Needed	25
System Is Needed to	Program Progress Difficult to Assess	26
Assess Program	DOE May Need to Manage ASCI as a Strategic System	28
Progress	Conclusions	29
	Recommendations	29
	Agency Comments and Our Evaluation	30
Chapter 3		31
Management and	Cost Estimates Increased Substantially	31
Tracking of Costs	Better Oversight of Costs Needed	35
Need to Be Improved	Conclusions	35
	Recommendation	36
	Agency Comments and Our Evaluation	36
Chapter 4		37
Technical Challenges	Technology Development	37
Are Present in All	Integration of Scientific Data	41
Aspects of the ASCI	Technical Expertise	43
Program		
Appendixes	Appendix I: Program- and Laboratory-Level Software Milestones From DOE Planning Documents, Fiscal Years 1996-2004	46
	Appendix II: Comments From the Department of Energy	49
Table	Table 1.1. Requested ASCI Funding by Program Component, Fiscal Year 1999	16

Figures

Figure 3.1 Original and Current Cost Estimates for the ASCI Program, Fiscal years 1996 Through 2004	32
Figure 3.2: Allocation of Estimated ASCI Costs by Major Program Areas, Fiscal Years 1996 Through 2004	33

Abbreviations

3-D	Three dimension
ASCI	Accelerated Strategic Computing Initiative
DARHT	Dual-Axis Radiograph Hydrodynamic Test Facility
DOE	Department of Energy
GAO	General Accounting Office
NIF	National Ignition Facility

Introduction

Historically, the United States used actual nuclear detonations as the primary method of validating designs and certifying the weapons as safe and reliable. Since September 1992, there has been a moratorium on testing. To ensure the continued safety and reliability of nuclear weapons, the Department of Energy (DOE), which is responsible for designing and building nuclear weapons, developed the 15-year Stockpile Stewardship and Management Program in 1995 as a substitute for actual testing. The stockpile stewardship program employs a variety of means to ensure weapons' safety and reliability, including examining weapons, conducting laboratory experiments and tests, and conducting computer modeling and simulation. The computer modeling and simulation part of the program is known as the Accelerated Strategic Computing Initiative (ASCI). The ASCI program aims to replace actual testing with advanced computer models that will simulate nuclear detonations. This effort requires modeling in 3-dimensions (3-D), with higher resolution than previous models and with better treatment of the underlying physical processes that occur during an actual nuclear detonation. To run the models, DOE is developing, as part of the ASCI program, the largest and fastest computers, which may ultimately be able to perform 100 trillion mathematical operations per second—10,000 times more powerful than those used to design the weapons originally. The ASCI program is expected to cost about \$5.2 billion for fiscal years 1996 through 2004.

The Evolving Role of Computing in Nuclear Weapons

Computers have been used to design and build nuclear weapons almost from the dawn of the nuclear era. As early as 1945, designers began using the ENIAC—the world's first computer, built at the University of Pennsylvania with government support—to perform calculations on the viability of a hydrogen or thermonuclear bomb. A successor version, which was fully electronic¹—the MANIAC—was built at Princeton in 1949, and a duplicate was built at Los Alamos. From that time, computers, and later so-called supercomputers, would play an increasing role in the designing and building of the U.S. stockpile of nuclear weapons. Computer models were used to design weapons and to interpret data from actual nuclear weapons tests. Models and computers were also used to identify and evaluate problems in the nuclear weapons stockpile. In the end, however, the final arbiter of a weapon's safety and reliability was usually an actual test or series of tests.

¹The ENIAC used vacuum tubes instead of gears to perform calculations but had to be programmed for each new problem by physically rearranging its circuit wires, which looked like old-fashioned telephone switchboard cords.

Since the first nuclear weapon test, known as Trinity, on July 16, 1945, the United States has conducted over 1,000 nuclear weapons tests. Testing was the principal method used to certify the safety and reliability of nuclear weapons. Testing was used to demonstrate that a particular weapon design actually worked and yielded the expected power and to prove the safety and reliability of components. For example, testing could be used to demonstrate that older components were still functioning properly after years of exposure to extremes of heat and cold and to radiation. In addition to periodically testing stockpiled weapons, the United States frequently developed new weapons to replace older weapons in the stockpile, thus ensuring the continued reliability and safety of its arsenal.

In September 1992, the Congress imposed a 9-month moratorium on underground nuclear testing.² This moratorium continued to be observed until September 1996, when President Clinton signed the Comprehensive Test Ban Treaty.³ The test ban treaty has been interpreted by the Administration to mean that no underground testing is allowed that results in any nuclear yield—no matter how low. The practical result of the test ban treaty is that existing nuclear weapons will be kept longer than planned because new weapon designs cannot be tested and certified as safe and reliable. The longer life span of the existing stockpile of nuclear weapons increases the possibility that they will decline in either performance or safety because of age-related factors like extended exposure to heat, vibration, and radiation. Faced with these testing restrictions, DOE developed a new approach to certifying the safety and reliability of weapons in the U.S. stockpile. A 1994 “Nuclear Posture Review” charged DOE with maintaining the capability to design, fabricate, and certify new weapons, if that ever became necessary. DOE responded by developing the 15-year Stockpile Stewardship and Management Program in 1995. The program is intended to ensure the continued safety and reliability of existing nuclear weapons using a variety of means, including examining weapons to find possible problems, conducting experiments to predict problems, and deciding on the basis of the results of these efforts what, if anything, needs to be done to ensure the continued reliability and safety of the weapons.

²Atmospheric testing was banned in 1963.

³Although the U.S. Senate has not yet ratified the treaty, a statutory extension of the 1992 moratorium took effect on September 30, 1996, and continues “unless a foreign state conducts a nuclear test” after that date, in which case the moratorium is lifted. DOE continues to observe the testing moratorium.

The ASCI component of the Stockpile Stewardship and Management Program was intended to provide the modeling and computers necessary to simulate in great detail the detonation of a nuclear weapon. Related experimental facilities like the National Ignition Facility (NIF) located at Lawrence Livermore National Laboratory, and the Dual-Axis Radiograph Hydrodynamic Test Facility (DARHT), located at Los Alamos National Laboratory, are intended to provide the data needed to address basic physics questions and to validate the accuracy of the ASCI computer models. With this change to a science-based rather than a physical test-based approach to addressing stockpile issues, the ASCI program has become a critical link in certifying the safety and reliability of nuclear weapons. The ASCI computer models and hardware will be used to identify potential stockpile problems by predicting the effects of aging and the need to replace components or even to retire weapons systems if they become unsafe or unreliable. In addition, the ASCI program will be used to design and certify needed replacement parts as well as the entire weapons system.⁴

The existing stockpile of weapons is aging, and many of the designers of those weapons have retired or are approaching retirement. For these reasons, DOE has decided that it is crucial to have the ASCI program available by fiscal year 2004, including the models and computers capable of performing 100-trillion operations per second. The intent is to have the remaining designers compare the output of the models against their actual experience with nuclear weapons tests as one means of validating the accuracy of computer models.

Why Such Large Computers Are Needed

The current generation of nuclear weapons were designed on computers that were much smaller than those being developed for the ASCI program—several million or a few billion operations per second versus 100-trillion operations per second. A logical question rises as to why such vastly larger computers are needed to ensure the safety and reliability of existing weapons compared with those computers that were needed to design and build these same weapons to the same safety and reliability standards.

The current stockpile of nuclear weapons were designed and built using much less capable computers and far simpler models than those

⁴Many of the manufacturing processes and technologies that were used to build the current generation of nuclear weapons and the components that they contain no longer exist. As such, replacement components manufactured using new processes, technologies, or materials need to be tested, in some manner, and certified as to their performance and impact on the weapons performance.

envisioned for the ASCI program. These less capable computers could have been used for several reasons, including (1) key components of the weapons were designed with a high level of symmetry so that a one- or two-dimensional view of the component would be fairly representative of the whole component, (2) weapons were designed without a need to model all of the underlying physics, (3) actual testing was used to resolve any uncertainties, and (4) weapons were routinely replaced by newer, tested weapons before they reached the end of their design life. With the loss of testing opportunities and the aging of the current stockpile, this approach is no longer feasible. Instead, DOE believes it is necessary to provide detailed visual 3-D simulations of nuclear weapons processes (that is, virtual testing capability).

Virtual testing requires far more complex and detailed models and much greater computer capability to run these highly complex models in a reasonable period of time. For example, to run certain two-dimensional weapons calculations on a Cray YMP supercomputer (an old generation of supercomputer but the type in use when some of the existing weapons were designed) took up to 500 hours. By comparison, moving from a two-dimensional to a 3-D model without changing any other parts of the model results in a calculation that is 1,000 times larger. At the same time, better detailed physics calculations of what is happening at the time of the nuclear detonation could require a calculation that is another 100,000 times larger. By extrapolating from these estimates, DOE concluded that running such a calculation in a reasonable amount of time (generally no more than several days for the largest calculations) would require computers capable of calculating at the rate of 100-trillion operations per second. Such machines were far beyond those commercially available when the ASCI program was started. Developing these increasingly powerful machines is one of the main goals of the ASCI program. Building the highly complex 3-D models is another.

ASCI Program Overview

The ASCI program is comprised of several components. As shown in table 1.1, the key components are Applications (software development), Platforms (computers), and Infrastructure (peripheral technologies such as networks, storage, and visualization). The program also includes the Academic Strategic Alliances Program, which contracts with universities for computing and scientific research to complement ASCI efforts. In fiscal year 1999, the scope of the ASCI program expanded when three new components were created. Numerical Environment for Weapons Simulation will acquire the infrastructure hardware needed for data

management and visualization. Distributed Distance Computing will provide the infrastructure needed to extend the advanced computing capabilities of the ASCI program to remote DOE weapons facilities. Verification and Validation is intended to verify that ASCI software is executing calculations as intended and to validate the accuracy of ASCI software results. Table 1.1 shows funding for ASCI-related activities, including existing computing facilities, verification and validation, and other activities at each of the three laboratories to support ongoing stockpile stewardship program requirements.

Table 1.1. Requested ASCI Funding by Program Component, Fiscal Year 1999

Dollars in millions	
ASCI program component	Requested funding
Applications (software)	\$152
Platforms (computers)	70
Infrastructure	
Problem Solving Environments	46
Numerical Environment for Weapons Simulation	31
Distributed Distance Computing	28
Verification and Validation	13
Stockpile Computing	156
Academic Strategic Alliances Program	14
One Program/Three Labs (program coordination)	6
Total	\$516

ASCI activities are carried out by DOE's three Defense Program laboratories—Los Alamos, Lawrence Livermore, and Sandia National Laboratories—with guidance from DOE's Office of Strategic Computing and Simulation under the Assistant Secretary for Defense Programs. For hardware development, the ASCI program intends to build on economically viable computer vendor technologies and thereby foster the health of the U.S. computing industry while also stimulating competition in this industry to adopt new technologies for advanced computing. ASCI computers of varying sizes will be built and housed at each of the laboratories, and their capabilities will be accessible to all three laboratories. Infrastructure-related hardware will be procured from available vendor technologies.

For software development, the ASCI program is relying on coordinated efforts at the three laboratories, supplemented by university-based

research when possible. At Los Alamos and Livermore, software development efforts focus on models that simulate the performance of the nuclear components of weapons systems. In contrast, software efforts at Sandia focus on models that simulate the performance of nonnuclear weapons components like the arming, firing, and guidance systems. Infrastructure-related software is being developed in a joint effort by the three laboratories.

The schedules for hardware, software, and infrastructure development are interdependent. For example, ASCI software must be able to operate on a variety of increasingly large parallel computer systems. The development of such “scalable” software requires the availability of computers and peripheral technologies that are sufficiently advanced to test and develop the software. Major milestones for hardware, software, and infrastructure development have been established for the program to 2004. Executing the ASCI program and meeting these milestones with the involvement of three laboratories will require close integration among programs and across laboratories.

Although the ASCI program’s ultimate goal is to provide 3-D weapons simulation capabilities by 2004, the ASCI computers and software developed to date are already important tools for addressing DOE’s high-priority stockpile needs. Today’s ASCI computational capabilities, for example, are being used to help design scientific experiments and to support the revalidation and certification of certain weapons and/or their components in a simulation environment.

Objectives, Scope, and Methodology

As requested by the Chairman, Subcommittee on Military Procurement, House Committee on Armed Services, we reviewed the management of the ASCI program, including (1) whether the program was meeting its key milestones and whether its hardware and software developments are adequate to date; (2) whether the program was within its projected budget; and (3) what key technical risks the program faces. The scope of this review encompassed all aspects of the ASCI program, which is conducted primarily by the Lawrence Livermore, Los Alamos, and Sandia National Laboratories.

To determine whether the program was meeting its key milestones, we obtained and reviewed planning and tracking documents and interviewed ASCI program officials from the three weapons laboratories and from DOE’s Office of Defense Programs. We visited each of the three laboratories and

DOE to discuss the program's progress in meeting its key milestones and to obtain evidence to verify the statements made by program officials. In July 1998 and January 1999, we attended the semiannual ASCI "principal investigator" meetings to learn more about the program's progress. We also reviewed studies and reports that have assessed the status and progress of the ASCI program.

To determine whether the program is within its projected budget, we examined cost and budget information provided to us by DOE and the three laboratories. We also reviewed information from DOE's Financial Information Variance Reporting System and contracts. Furthermore, we examined budget information that was included in DOE's budget request for fiscal years 1996 through 2000. We spoke to program officials to determine how costs were estimated, why program costs have escalated, and how they review and manage laboratory costs. We did not independently verify the reliability of information contained in DOE's financial management system, which we used in this report, because it is the basis for DOE's financial statements, to which we have given an "unqualified opinion" in our audit of the federal government's financial statement.

To identify the key technical risks facing the program, we obtained and reviewed program-planning documents and interviewed ASCI program managers from the three weapons laboratories and from DOE's Office of Defense Programs. We also reviewed studies and reports on the ASCI program and other materials related to high-performance computing.

We conducted our review from July 1998 through June 1999 in accordance with generally accepted government auditing standards. Key contributors to this report were Linda Chu, Daniel Feehan, Anne McCaffrey, and Edward Zadjura.

A Comprehensive Planning and Tracking System Is Needed to Assess Program Progress

Weak management and information processes hamper oversight of the strategic computing program and make assessing progress towards program milestones difficult and subjective. Although program managers report that many milestones have been met, the lack of comprehensive planning and progress tracking systems make an assessment of the program's short- and long-term progress difficult and subjective. Current planning efforts include a strategic plan that is out of date, annual plans that have been prepared sporadically, and milestones that are not well defined. Efforts to track the program's progress are not consistent, and no clear record exists of program accomplishments compared with milestones. Consequently, it is difficult to determine which of the hundreds of milestones have been met, which are behind schedule, or even which are still relevant, given changes in program priorities, and how progress on individual projects contributes to the program's overall goals. In response to our requests for information, ASCI program officials have begun to institute more systematic procedures for planning and for tracking and reporting program progress.

Comprehensive Planning Is Needed

The ASCI program's long- and short-term planning efforts thus far have been inconsistent and incomplete. Strategic planning documents have not been updated since the program's inception. During the program's first 3 years, annual implementation plans were prepared inconsistently or, in some cases, not at all, resulting in an incomplete program baseline. The long-term ASCI milestones presented in various plans are inconsistent, and no information exists to link annual activities to these milestones. Strategic plans also do not identify the multiple research strategies currently employed to meet many long-term milestones, nor do the plans include decision points for managing these strategies. Performance criteria for most milestones have not been defined in the planning process.

Strategic and Annual Plans Are Outdated, Inconsistent, and Incomplete

The plans used to manage the ASCI program to date have numerous limitations. The ASCI Program Plan is the program's primary strategic plan. Published in 1996, the plan included a list of long-term (program-level) milestones for hardware, software, and infrastructure development. Although hardware and infrastructure milestones have not changed much, program officials have revised software milestones numerous times but have not published an updated strategic plan. Despite a proliferation of program planning documents showing software milestones, there is little consistency among these documents, and no clear record of when and why milestones were changed and which are the most current. Some of

the original milestones are intact or have changed little over time, others have changed considerably, and still others have been replaced with new milestones. In addition, the dates for certain milestones are inconsistent among documents. For example, the “Prototype 3-D Primary Simulations” milestone for fiscal year 1999 has been shown consistently, although the milestone name has changed slightly. By contrast, several milestones do not appear consistently in different documents. For example, the “abnormal environment thermal assessment” milestone is listed as occurring in fiscal year 1999 in one planning document and in fiscal year 2000 in another. (App. I lists ASCI program milestones as identified in various planning documents).

This lack of consistency is also found in another strategic planning document—the Stockpile Stewardship Plan (the “Green Book”). The Green Book is the strategic plan for DOE’s Office of Defense Programs and is supported by more detailed planning documents, including the ASCI Program Plan. The ASCI milestones presented in the Green Book, however, are not always consistent with those included in ASCI program documents. For example, the Green Book lists the “crash/fire safety” and the “full physics, full system prototype” as milestones, although they are not listed among the most critical milestones compiled by the program office. Conversely, the program office’s most critical milestones list includes the “hostile environment electronics assessment” and the “abnormal environment thermal assessment” milestones, but the Green Book does not list either. Program officials explained that the Green Book is prepared at different times and for different purposes than ASCI planning documents. Nonetheless, ASCI milestones in DOE’s Green Book and the ASCI Program Plan should be consistent.

In another example of inconsistent long-term planning, the Simulation Development Roadmap has never been updated. This document was intended to ensure that the ASCI program’s simulation and modeling activities would be clearly identified and that priorities would be set on the basis of stockpile needs and current capabilities. DOE intended to update this plan periodically as the program progressed. Neither DOE nor the laboratories, however, have developed this document beyond the initial planning level since 1996. No other equivalent source of information exists that provides a long-term, needs-based perspective on the ASCI program’s modeling and simulation activities. DOE program officials said that there is no programmatic requirement to update the Simulation Development Roadmap and that some information about program needs is contained in annual plans. This statement contradicts language in the

Stockpile Stewardship Green Book, which states that the Simulation Development Roadmap is to be used in conjunction with annual planning documents to define computing requirements. Furthermore, as discussed below, annual planning documents have not been consistently prepared.

The three laboratories in this program have also developed software milestones (laboratory-level milestones) and published various compilations of these. A consolidated set of these milestones, known as the “Consolidated Roadmap” is supposed to show the software milestones for all three laboratories. Another listing, known as the “Nuclear Roadmap,” shows combined software milestones for two laboratories—Los Alamos and Livermore—whose modeling efforts focus on the nuclear components of weapons systems. A third listing, known as the “Non-Nuclear Roadmap,” shows software milestones for Sandia Laboratories, whose modeling efforts focus on nonnuclear weapons components, such as those for arming, firing, and guidance. As with other ASCI program documents, these sources report milestones inconsistently, and laboratory-level milestones may or may not match program-level milestones. For example, the consolidated roadmap lists milestones such as “3-D forging/welding microstructure,” “full physics burn code prototype,” and “burn code with aging” that do not appear in either the nuclear or nonnuclear roadmaps. Conversely, the nuclear and nonnuclear roadmaps list milestones such as the “3-D nuclear safety simulation” and the “B61 penetrator” that do not appear in the consolidated roadmap.

The laboratories’ long-term planning efforts also include multiyear plans for some of the individual software projects that contribute to laboratory-level milestones, but such plans are not required, and their format has varied from laboratory to laboratory. At Sandia, for example, project plans spanning 5 years have been prepared for key software projects. At Los Alamos and Livermore, project plans also have been developed for key software projects, but their time frames are shorter and variable (that is, 2 or 3 years).

Short-term planning, as represented by annual implementation plans, has also been inconsistent and incomplete. The purpose of these plans is to specify project tasks and milestones for the current year (annual milestones). However, during the first 3 years of the program, annual implementation plans were prepared by the three laboratories for some, but not all, components of the ASCI program and for only some of those years. DOE’s failure to ensure comprehensive and consistent planning during those years has resulted in an incomplete program baseline for

fiscal years 1996 through 1998. In addition, the annual plans do not explain changes in the scope of the work or milestones that occur from year to year. Fiscal year 1999 was the first year that DOE required the laboratories to submit implementation plans for all the components of the program and consolidated these documents into a program-wide ASCI Implementation Plan.

The Linkages Between Annual and Long-Term Milestones Need to Be Identified

To date, DOE has not required that program documents show linkages between annual milestones and long-term program- or laboratory-level milestones. With the exception of information provided in response to our request, these linkages have never been documented. Without clear information to identify such links, it is impossible to determine how annual progress contributes to meeting those long-term milestones. Program officials maintain that such links do exist but that they are not explicit in program documents. They also acknowledged that technical expertise is needed to identify links between annual milestones and long-term milestones. In our discussions with ASCI laboratory staff, however, we found that such links do not always exist and sometimes could not be identified even by laboratory personnel. At Sandia, for example, laboratory officials identified ASCI software activities that are needed to meet stockpile requirements but that are not tied directly to program-level milestones. In another case, it was not apparent what laboratory activities contributed to meeting the “macro-micro aging” milestone. DOE and laboratory officials identified Sandia as the laboratory responsible for this milestone, planned for completion in early 1999. While Sandia officials identified some activities that they believed were relevant to meeting this milestone, they were unsure about whether Sandia’s activities were all that was needed to meet this milestone. They said it was possible that ongoing projects at the other two laboratories contributed to meeting this milestone.

Program and laboratory officials agreed that such links should be made more apparent and, in trying to respond to our request for information about these links, they attempted to identify and document linkages. Although complete information was not provided by all of the laboratories, the information received shows that annual milestones are not always directly linked to long-term milestones. DOE and laboratory managers also told us that the process of developing this information was helpful for tracking the progress of the program and that they plan to refine this process and update the information on a quarterly basis.

Plans Should Clearly Identify Research Strategies, Critical Paths, and Decision Points

ASCI plans do not identify all of their multiple research strategies nor do they establish decision points for identifying which of these strategies are critical for meeting key program milestones. Because of the complex technological challenges involved in developing software and hardware to model nuclear weapons, the laboratories have undertaken multiple research strategies in an effort to mitigate risk and achieve their laboratory- and program-level milestones. According to laboratory officials, they will eventually have to reassess these strategies to see which are working and which are not and decide how they can focus their efforts to best achieve their milestones. As a result, individual projects could be enhanced, scaled back, or eliminated as “critical paths” are defined.

In the area of infrastructure development, for example, one overall goal is to develop ways for scientists to examine massive amounts of weapons simulation data. To achieve this goal, the development effort has at least seven simultaneous lines of effort, including developing “common data formats” and “distributed file systems.” Within those lines of effort, multiple research approaches are being pursued. A program official explained that all of these lines of effort are needed but that it is not clear at this time which efforts will be critical to meeting the overall goal. It is also not known when a choice among these options needs to be made. Another development effort is focused on specialized software that would help ASCI software developers understand and improve the performance of their weapons software programs, which ultimately will help to reduce the time needed to solve such problems. Currently, there is no such specialized software to solve such problems. Commercial vendors are developing such software, but only to a certain extent because there is limited demand for such products outside the ASCI program. As a result, while ASCI program officials are considering vendor products, the program is also funding development efforts at universities and collaborating with industry to develop the needed software. The program official explained that the software needed for the ASCI program may be available from at least one of these sources, but, again no decision points have been established for when a choice among these options needs to be made. According to this official, infrastructure research strategies are managed and decisions about them are made on a yearly basis by a team assembled from the three laboratories.

In the area of software development, ASCI software developers at Los Alamos and Livermore laboratories use multiple research and risk mitigation strategies in developing their weapons simulation software. Los Alamos and Livermore have multiple software teams competing to develop

weapons simulation software. Program officials explain that having several teams is advantageous because the arrangement allows the laboratories to explore different simulation approaches, cross-check/validate each other's work, and mitigate the risk of selecting one approach before all approaches are explored adequately. These different research and risk mitigation efforts, however, are not identified clearly in planning documents. In addition, although laboratory officials acknowledge that they will need to eventually scale back or restructure their software efforts, no decision points have been established for doing this.

Plans Should Define Performance Criteria for Milestones

Measurable performance criteria for most ASCI milestones have not been defined. Laboratory officials were, in many cases, unable to specify what needed to be done to complete a milestone, which laboratory(ies) bore responsibility for meeting it, or what their own particular contributions to meeting that milestone were. Laboratory officials said that no objective and specific measures exist, in most cases, to determine whether milestones have been successfully completed. Instead, they make subjective judgments about when a body of work meets a given milestone. For example, the "micro-aging" milestone, which appears in plans as a critical program milestone, was scheduled for completion in 1997. Los Alamos officials said they were unfamiliar with this milestone but believed it was Sandia's responsibility. Sandia officials agreed that they were responsible for this milestone and said that, in their opinion, the milestone had been met. However, while they identified relevant tasks, they were unable to identify precisely and completely what was supposed to be done to meet the milestone. Sandia officials speculated that the other two laboratories also had a responsibility for meeting this milestone.

On the other hand, we found one example of a milestone for which specific performance criteria have been established. The "Prototype 3-D Primary Simulations" milestone, scheduled for completion by the end of calendar year 1999, was subject to a review by program officials in June 1998. The review established technical specifications for successfully completing the milestone and assessed the progress of the Los Alamos and Livermore software development activities contributing to this milestone. A progress review for this milestone is planned for the summer of 1999 and a follow-up review is planned upon its completion. Program officials acknowledge that they need to perform similar reviews for other milestones. In May 1999, DOE issued a report on its review, entitled "Codes for the Complex," of the nonnuclear mechanics software development

efforts. At this juncture, however, most program-level milestones remain undefined.

Systematic Progress Tracking and Reporting Needed

The efforts of DOE and laboratory managers to track ASCI progress thus far have been inadequate. Progress tracking has been limited primarily to the reporting of annual accomplishments without any systematic tracking of progress towards long-term milestones. Current tracking efforts occur through a variety of formal and informal methods, and the format and organization of these tracking efforts has varied from year to year. As a result, it is not possible to determine whether annual milestones were achieved or the extent to which annual efforts were contributing to the laboratory- and program-level milestones. In response to our request for tracking information, program officials decided to track and report the program's progress more systematically.

The ASCI Program Plan described the semiannual principal investigators meetings as the primary forum for reporting program progress. The plan also noted that performance metrics would be developed and used at these meetings to compare actual output with planned output. However, these metrics were never developed. While these meetings are a forum for ASCI researchers to exchange ideas, there are few reporting requirements, and reports on program accomplishments have generally not been related to established milestones. Furthermore, the meetings focus on presentations of individual projects, with no effort to pull together a systematic and comprehensive assessment of how the ASCI program is progressing towards its overall goals.

Laboratory officials met with DOE officials in fiscal years 1997 and 1998 to report their annual accomplishments. These meetings were focused solely on accomplishments and not on reporting the status of all milestones. As a result of this limited reporting process, some key information about the program's progress was not divulged. For example, accomplishments reported for fiscal year 1998 for one key software project did not reveal that the project was actually 6 months behind schedule, which affected the schedule of other related projects.

The Annual Performance Report is published at the end of each fiscal year to report on the progress of the Stockpile Stewardship Program, including the ASCI efforts. The report discusses various ASCI milestones met during the year, but these accomplishments are not all tied to the program-level milestones, and no assessment is made of how the program is progressing

in terms of meeting its most critical milestones. As such, the report did not provide a comprehensive assessment of the ASCI program. However, according to the Green Book, the ASCI program is critical to the success of the overall Stockpile Stewardship Program.

To varying degrees, the laboratories used the annual implementation plans to report accomplishments during fiscal years 1996 through 1998, although this practice was discontinued in fiscal year 1999. In fiscal year 1997, the implementation plan for infrastructure work at the three laboratories did not report any accomplishments for 1996, while the Sandia plan for software development reported several. However, the reported accomplishments cannot always be correlated easily with established milestones. In some cases, a connection was apparent, while in other cases it was not. In response to our request for information, the laboratories agreed to prepare material showing progress in meeting the milestones established in recent annual implementation plans. Because implementation plans were not prepared for all ASCI program components in 1998, however, draft documents and other plans are being used as a baseline to track progress in certain areas. Laboratory and DOE officials said that the information they developed at our request was very helpful for tracking program progress. DOE plans to have the laboratories refine their efforts and require that this information be updated on a quarterly basis.

Program Progress Difficult to Assess

Although program managers have reported that many milestones have been met, it is difficult to gauge the ASCI program's overall progress because of weaknesses in program management and information processes. The lack of a systematic progress tracking system, combined with the lack of defined performance criteria, make it difficult to assess whether the ASCI program is proceeding on schedule and delivering the performance expected. However, it is possible to gain some limited insights by discussing individual projects or areas with laboratory officials. In an attempt to determine the adequacy of hardware and software development to date, we discussed program progress in these areas with program officials. The insights we gained are discussed below.

Hardware Performance

In the area of hardware development, most contract milestones to date that relate to the delivery and installation of the computers and related hardware have been met, although not all acceptance tests have been met. Currently the two 3-trillion operations per second systems at Los Alamos

and Livermore are operational. However, Los Alamos is still working to address problems affecting the reliability and stability of its system, which, as a result of these problems, did not pass all its contractually required acceptance tests on time.

Progress in hardware development has been reported in annual implementation plans, at semiannual principal investigator meetings, and in press releases to the media. DOE and the laboratories, for example, have issued press releases that emphasize how fast these systems have performed on tests of their sustained speed. These releases may be misleading in some cases because they are reporting on only selected aspects of performance. For example, DOE reported that the Los Alamos system established a world record for sustained speed. However, it was not reported that the test used was substantially easier than the test specified in the contract or that the system has not yet met other acceptance test criteria (such as, mean time between failures). The sustained speed of this machine would be one-half the speed reported in press releases had the test specified in the contract been used.

In addition, the Livermore computer continues to have problems with parallel input/output file operations. DOE's high-performance computers are expected to incorporate state-of-the-art hardware and software technologies. These computers can process multiple parts of one program at the same time, using parallel-processing techniques. According to the ASCI Program Plan, of the 100,000-fold increase in computing performance needed by the program, DOE expects a 10-fold increase from improving software to take advantage of parallel-processing techniques. However, the management of the input and output data during such processing continues to be an issue. For example, the data that result from running a model on 1,000 processors currently have to be saved to 1,000 separate files, making the data more difficult to manage and use. The laboratory is working with the vendor to address this problem.

Software Development

Documentation of ASCI's software performance is limited. With the exception of the "Prototype 3-D Primary Simulations" milestone review discussed previously, little documentation exists that compares software development progress against established milestones. This review, also known as the "burn code review," is the only formal review of a milestone done to date. The numerous other software projects related to other milestones, such as 3-D secondary burn code projects at Los Alamos and Livermore and all software projects at Sandia, have not been reviewed.

According to our analysis of planning documents, accomplishment reports, and other related material on software activities at Sandia, the ASCI program's current planning and tracking system does not always provide an accurate picture of progress. For example, in reviewing the documentation for a major software development effort that Sandia officials told us was 6 months behind schedule, we found no indication that delays had occurred. The documentation showed that most fiscal year 1998 milestones had been met, and all milestones for the first quarter of fiscal year 1999 also were met. Sandia officials explained that, although the reported information was accurate, milestones had not been established for all of the important tasks associated with that project. Since the purpose of this project is to provide a common framework for all Sandia ASCI software models, delays in this project are hindering the performance of other software projects that require integration into the common framework. Sandia also did not report these delays in the list of accomplishments for fiscal year 1998 that it submitted to DOE. Sandia officials told us that, as of March 1999, 50 percent of their software development projects were experiencing delays because of funding or program changes or are behind schedule for other reasons. In general, detailed performance requirements have not been established for most software milestones, so it is difficult to develop an objective assessment of performance in this area. In effect, the judgment of whether software tests or demonstrations have achieved the desired level of performance is based on the subjective opinion of the laboratory and DOE program managers.

DOE May Need to Manage ASCI as a Strategic System

DOE may not be appropriately managing the ASCI program by not designating it as a strategic system. DOE has established criteria for designating its most important projects as strategic systems to ensure oversight at the highest departmental level. The criteria are that the project costs more than \$400 million, is an urgent national priority, be high-risk, have international implications, or be vital to national security. The purpose of designating strategic systems is to ensure informed, objective, and well-documented decisions for key events, such as changes to baseline cost or schedule. The ASCI program meets the criteria for being treated as a strategic system. The ASCI program will likely cost about \$5.2 billion for fiscal years 1996 through 2004, is a critical part of the stockpile stewardship program, is an urgent national priority on national security grounds, and has international implications because it is a major factor in U.S. support of the Comprehensive Test Ban Treaty. Finally, the ASCI program is high risk because it seeks to advance the state of the art in computers, modeling, and simulation well beyond current capabilities, has

already experienced delays, has had its projected cost increase, and depends on as-yet unknown technologies for success. These characteristics, coupled with the demonstrated weaknesses in program management and oversight, make the ASCI program a clear candidate for being designated as a strategic system. According to DOE, it has not designated this effort as a strategic system because the program is already subject to high-level department oversight. However, as discussed above, we found serious weaknesses in the program's management and information processes that make it difficult to determine if the program is performing as expected.

Conclusions

DOE's oversight of the ASCI program is hampered by weaknesses in management and information processes. The program lacks a comprehensive planning system—one that clearly establishes milestones; links short- and long-term milestones; identifies research strategies, critical paths, and decision points; and defines performance criteria for milestones. Furthermore, the program lacks a progress tracking and reporting system. Consequently, overall program progress is difficult to assess. DOE has chosen not to designate the ASCI program as a strategic system. The demonstrated weaknesses in the ASCI program's management and information processes, coupled with the program's critical role in DOE's mission to maintain the nation's stockpile of nuclear weapons without testing, warrant DOE's designating ASCI as a strategic system requiring the highest levels of management attention.

Recommendations

Given the ASCI program's lack of a comprehensive planning, tracking, and reporting system and the importance of the program to maintaining the stockpile of nuclear weapons, it is important that DOE improve its oversight and management of this program. Therefore, we recommend that the Secretary of Energy require the establishment of a comprehensive planning, tracking, and reporting system. This system should, at a minimum, establish clear milestones; identify links between short- and long-term milestones; identify research strategies, critical paths, and decision points; define performance criteria for the successful completion of milestones; and establish progress tracking and reporting requirements. We further recommend that the Secretary of Energy designate the ASCI program as a strategic system warranting oversight at the highest departmental level.

Agency Comments and Our Evaluation

DOE agreed with the recommendation to improve the oversight and management of the program and cited several changes that it was making. Specifically, DOE stated that it would soon issue an updated Program Plan that will include detailed specifications for all of the critical program milestones. In addition, the FY 2000 Implementation Plans that will be issued by September 30, 1999, will include descriptions of all program elements and complete lists of all milestones. The Department also cited the creation of a quarterly progress tracking mechanism to track program milestones. However, in addition to tracking the program's progress against established calendar milestones, it is necessary to establish specific technical criteria for what constitutes the successful completion of those milestones. Until DOE completes and publishes its revised Program Plan and FY 2000 Implementation Plans, we cannot determine whether it has fully complied with this recommendation.

DOE disagreed with the recommendation to designate the ASCI program as a strategic system and stated that to do so would duplicate the planning, progress tracking, and reporting system. We agree that creating a duplicative tracking system that mirrors the requirements set out by DOE for strategic systems would not be worthwhile. However, as discussed in detail in this report, DOE has not shown that it has an adequate planning, progress tracking, and reporting system in place for the strategic computing initiative. While DOE is making improvements in these areas, the changes are not yet fully in place and their adequacy cannot be judged at this time. Furthermore, if the changes that DOE is making are adequate to meet the requirements for tracking and monitoring a strategic system, then we cannot understand DOE's reluctance to designate this large and costly program as a strategic system. DOE stated that it has a review process that meets the intent of the Clinger-Cohen Act of 1996. However, we reported in July 1998 that the Department's process effectively excludes scientific computers like those being acquired through the ASCI program from DOE's normal review channels and places them within the program offices.⁵ We stated that all computers should be included as part of the normal DOE Clinger-Cohen review process.

⁵Information Technology: Department of Energy Does Not Effectively Manage Its Supercomputers (GAO/RCED-98-208, Jul. 17, 1998).

Management and Tracking of Costs Need to Be Improved

ASCI cost estimates have increased substantially. In 1995, DOE estimated that program costs for fiscal years 1996 through 2001 would be \$1.7 billion.¹ By 1999, estimated costs for those years increased to \$2.9 billion. DOE currently estimates that the program will cost about \$5.2 billion for fiscal years 1996 through 2004. Some of the cost increases result from changing program requirements. For example, in 1996, the United States shifted from a program based on nuclear testing to one based on computer simulations of weapon performance. For the ASCI program, the shift to computer simulations resulted in higher costs to acquire the latest and fastest computers and to develop advanced simulation and modeling software. The cost increases also reflect weaknesses in DOE's cost estimation and management. For example, DOE has difficulty determining technical requirements and then reliably estimating costs for state-of-the-art computers and software. In addition, DOE limits its cost oversight to reviews of aggregate laboratory spending and consequently cannot determine if the costs of specific projects at the laboratories are over or under budget.


Cost Estimates Increased Substantially

DOE's cost estimates for the ASCI program have increased substantially since 1995, when early budget projections were made. Costs for fiscal years 1996 through 2001 have increased from an original estimate of \$1.7 billion to the current \$2.9 billion. DOE's actual fiscal year 2000 budget request for ASCI, which totals \$692 million, is more than double the original fiscal year 2000 estimate made in 1995. Figure 3.1 shows for each fiscal year the original and current budget estimates.

¹The estimated cost figures provided in this report have not been adjusted to constant dollars. Rather, they reflect DOE's budgeting and planning process estimates, which were provided in current dollars.

**Chapter 3
Management and Tracking of Costs Need to
Be Improved**

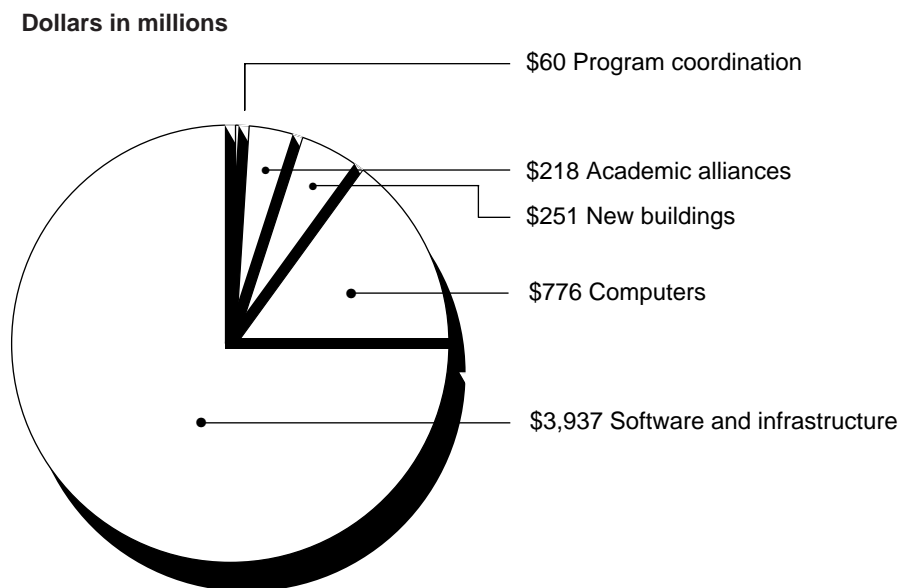
Figure 3.1 Original and Current Cost Estimates for the ASCI Program, Fiscal Years 1996 Through 2004



Source: GAO's analysis of DOE's data.

The current total estimated cost of the ASCI program, for fiscal years 1996 through 2004, is about \$5.2 billion. Although the program is scheduled to operate through 2010, estimates beyond 2004 have not been made. Figure 3.2 shows how this \$5.2 billion is allocated by program areas.

Figure 3.2: Allocation of Estimated ASCI Costs by Major Program Areas, Fiscal Years 1996 Through 2004



Source: GAO's analysis of DOE's data

The \$776 million estimated for computers includes the costs of the three computers currently installed at the laboratories and the planned acquisition of computers that perform 10-, 30-, and 100-trillion operations per second. The vast majority of ASCI funds, almost \$4 billion, will be spent on the software development and program infrastructure. Software and infrastructure include the development of the 3-D simulation models and the multiple peripheral technologies needed for, among other things, visualization, networking, and data management. DOE estimates it will spend \$251 million to construct new buildings at the laboratories that will house the computers, as well as offices and visualization theaters. The Academic Strategic Alliances Program with the universities will cost \$218 million, and another \$60 million is estimated for the coordinating laboratories' efforts.

Changing Requirements Account for Some Cost Increases

DOE officials told us that ASCI costs have escalated since the original budget projections because the program was originally intended as a limited effort to improve the computer capability available for stockpile

stewardship. In the original budget estimate, developed in 1995, DOE officials requested funding for a 1,000-fold increase in computing capability. Since then, the ASCI program has expanded because of changes in the U.S. nuclear weapons policy, particularly, the U.S. decision in August 1995 to pursue a “zero yield” Comprehensive Test Ban Treaty. This policy change meant that the United States would need to maintain the nuclear stockpile far beyond its design life and would have to shift from a traditional nuclear test-based program to one based on computer simulations to ensure the safety and reliability of nuclear weapons. As a result, DOE has developed strategies that focus on advanced modeling and simulation that require a 100,000-fold increase in computer capability.

Unreliable Cost Estimates Also Contribute to Cost Increases

While the change in program requirements has affected DOE’s budget estimates, officials also acknowledge their difficulty in estimating costs because of the unprecedented scale of the hardware and software technologies needed by the ASCI program. For example, before DOE began this effort, a computer with thousands of processors operating as a fully integrated system had never been built. In addition, software to run on systems of this size and high-performance visualization technology to display the results of simulations at this scale had never been developed.

According to one DOE official, the Department might lack the expertise to anticipate future technical requirements for state-of-the-art hardware and software and to reliably estimate their costs. For example, when planning the current, expanded ASCI program, DOE envisioned a single computer capable of 3 trillion operations per second that could provide access and collaborative opportunities to all three laboratories, using secure, high-speed networking capabilities. However, DOE decided to procure a second such computer in 1997 because the technical capabilities to support such long-distance computing were not yet in place. Combined, the two computer systems cost DOE almost \$220 million. In addition, DOE acknowledged that during early funding strategies, it did not consider the difficulty and importance of the technology needed by weapons designers to visualize the results of the 3-D weapons simulations. Yet such visualization technologies are required to graphically represent to weapons designers the results of 3-D ASCI simulations. DOE currently estimates that more than \$87 million is needed for visualization activities for fiscal years 1999 through 2004.

Better Oversight of Costs Needed

DOE's oversight of costs is limited to a review of aggregate spending at the laboratories. While DOE contends that cost controls for ASCI are in place, it does not track costs to determine which specific projects may be over or under budget. DOE monitors monthly spending for each laboratory but does not compare previously estimated costs for major projects with their actual costs. For example, as part of the budget formulation process, project costs are estimated and subsequently included as part of the Department's fiscal year budget request. However, DOE tracks only how much the laboratories have spent in broad categories that lump together costs for many projects. As a result, DOE cannot determine which projects, if any, may be costing more or less than originally planned. DOE told us it relies on the laboratories to determine whether projects are within their planned budget, but one laboratory ASCI manager told us that the laboratory tracks only the technical status of projects, not their costs.

In addition, DOE's limited oversight of the laboratories' activities could result in DOE's underestimating ASCI program costs. DOE estimates that ASCI program costs have increased by about \$1.2 billion compared with its original estimate for fiscal years 1996 through 2001. However, delays in completing projects at the laboratories could increase those costs. For example, as noted in chapter 2, one laboratory estimates that 50 percent of its software development projects are experiencing delays or are behind schedule. DOE's lack of information about the progress of projects, combined with its limited cost tracking, do not allow DOE to determine how much longer it will take to complete those projects or at what cost, thus limiting its ability to accurately project ASCI program costs.

In January 1999, we reported on the significant management challenges at DOE, including the difficulty completing large projects within budget.² We noted that DOE often requires large projects costing hundreds of millions of dollars that are often the first of their kind and involve substantial risk. ASCI is such a challenge. ASCI is critical to DOE's mission, is estimated to cost about \$5.2 billion, requires the development of hardware and software on an unprecedented scale, and involves substantial risks.

Conclusions

ASCI costs have increased substantially because of changes in program requirements and weaknesses in DOE's cost estimates. Because its tracking of costs is limited, DOE cannot determine whether specific projects are under or over budget. Historically, DOE has had difficulty managing the costs of large programs.

²Department of Energy: Major Management Challenges and Program Risks (GAO/OGC-99-6, Jan. 1999).

Recommendation

Given the substantial increases in the ASCI program's cost estimates to date, DOE's acknowledged problem in estimating costs for the unprecedented scale of development efforts involved in the ASCI program, and the lack of a cost-tracking process, it is important that DOE improve its oversight of ASCI program costs. Therefore, we recommend that the Secretary of Energy require that ASCI adopt systematic cost-tracking procedures that will allow DOE managers to determine if specific projects are within budget.

Agency Comments and Our Evaluation

DOE did not agree with our recommendation to adopt systematic cost-tracking procedures for the strategic computing initiative, noting that funding and costs are tracked by budget and reporting codes in the Department's Financial Information System. DOE stated that these systems are extended to individual projects using other funding and cost-monitoring tools that gather more detailed information. As an example, DOE cited an analysis performed in March 1999 of selected projects that identified the commitments and cost status for specific procurements at the project level. We do not agree that DOE has an adequate system for tracking at the project level or that the changes it is making will rectify this problem. DOE's current system tracks cost only at the aggregate level and does not allow DOE managers to determine which projects at the laboratories are under or over budget. Furthermore, the "other funding and cost monitoring tools" that DOE uses do not allow the systematic tracking of project costs. DOE also stated that some budgeting flexibility is necessary to capitalize on changes within the high-computing industry. While we agree that some budgeting flexibility is necessary in a project of this size and complexity, we do not believe that that flexibility should preclude the effective oversight of a multiyear program costing over \$5 billion.

Technical Challenges Are Present in All Aspects of the ASCI Program

The primary challenge facing the ASCI program is to develop a simulation capability that, in the absence of nuclear testing, can be used to determine whether a modified weapon system will perform as intended. The need for this “virtual test” capability encompasses most of the technical challenges associated with the ASCI program. These challenges range from developing state-of-the-art hardware and software technologies, to the integration of scientific data from weapons physics experiments, to recruiting and retaining staff with the technical expertise needed. The risks in failing to address these challenges are intensified because the program is a research-and-development effort with an accelerated schedule. Program officials acknowledge the multitude of risks associated with the program and point to the risk mitigation strategies they have designed to address program risks. These strategies include the use of several vendors to develop computers, overlapping software development efforts, and partnerships with industry and academia.

Technology Development

The development of hardware and software technologies and of the necessary infrastructure to support these technologies is critical to achieving the ASCI program’s simulation and modeling goals. Hardware development must successfully increase computational speeds to 100-trillion operations per second in 2004. Software development efforts are extensive and must ultimately incorporate massive amounts of data, solve progressively more difficult problems, and be capable of running on increasingly larger and faster computers. Developing the needed infrastructure, including data storage and visualization technologies, will require significant improvements.

Computer Speed Must Increase Dramatically by 2004

Developing computers capable of processing complex 3-D nuclear weapons simulations is one of the primary challenges facing the ASCI program. Increasing the computational speed to 100-trillion operations per second by 2004, according to program officials, is essential to meeting program goals. At Livermore, development is under way on a 10-trillion operations per second computer that is scheduled to be installed during fiscal year 2000, and the acquisition of a 30-trillion operations per second computer at Los Alamos is planned for fiscal year 2001. A request for proposals for the 30-trillion operations per second computer was sent out in May 1999, and contracts have been signed with several vendors to work on the related technology needed for a computer of this size.

Program officials explained that their risk mitigation strategy includes using competing computer vendors to independently develop increasingly larger computers at the three weapons laboratories. A 1.8-trillion operations per second computer that was developed by the Intel Corporation is in use at Sandia, while both Los Alamos and Livermore are developing computers capable of 3-trillion operations per second. The computer at Los Alamos is being developed by Silicon Graphics, Incorporated, while Livermore is working with the IBM Corporation to develop its computer. According to laboratory officials, the experience gained at Sandia, plus the competing efforts at Los Alamos and Livermore, helps to ensure that at least one of these computers will be generally available to carry out computational work. According to a program official, the competing computers at Los Alamos and Livermore are based on different technologies, which helps to further mitigate risks. Currently, Livermore is developing the 10-trillion operations per second computer with the IBM Corporation, while Los Alamos sent out a request for proposal to select a vendor to develop the 30-trillion operations per second computer in May 1999.

ASCI program officials explained that the ASCI computers being constructed involve thousands of processors, switches, disks, and related components that must work together as a fully integrated system to run the largest simulations. These officials explained that getting computer systems of this size to operate as a fully integrated system has never before been achieved and is one of the most difficult challenges facing the program. An April 1998 review of the computing division at Los Alamos by an external committee recognized this issue by pointing out that users generally had access only to small parts of the computer and rarely had access to the full system.¹ The Committee's report explained that operating the computer as a fully integrated system was important because the ASCI computer needs are based on running simulations that require the full capability of the computer. The challenge continues today with the 3-trillion operations per second computer at Los Alamos, which has experienced many failures when trying to run as a fully integrated system.

A March 1999 review by the ASCI Blue Ribbon Panel noted another important risk in meeting the schedules for computers operating in the range of 30- to 100-trillion operations per second.² The report explained that to meet the schedule for these larger computers, it might be necessary for the laboratories to write the system software necessary to enable the

¹Los Alamos CIC Division External Review Committee Report for the April 1998 Review.

²Report of the ASCI Blue Ribbon Panel, Mar. 2, 1999.

computer to operate as a fully integrated system. The report noted that there is a risk in such a course of action because laboratory personnel do not have extensive experience in this area.

**Software Development Is
Critical to Program
Success**

Developing software that incorporates all of the required science to simulate nuclear weapons while running on computers consisting of thousands of processors is, according to ASCI program managers, one of the most demanding tasks of the ASCI program. These officials explained that developing such software has historically taken approximately 5 years before it can be used with confidence. Because of the complexity, these officials stated that ASCI software may take longer to develop, and a key program goal is to reduce the development time to the 5-year historic average. A report by the President's Information Technology Advisory Committee described current software development as "among the most complex of human-engineered structures."³ The report noted that the nation's ability to construct needed software systems and to analyze and predict the performance of these systems is painfully inadequate. Part of DOE's risk mitigation strategy in ASCI software development includes competing software efforts at Los Alamos and Livermore. In addition, to leverage their efforts in software development, the laboratories have contracted with several universities (through the ASCI Academic Strategic Alliances Program) to conduct research in areas of high-performance computing and physical science.

The technical challenges inherent in the development of ASCI-related software are due in part to the complexity of the needed software. Program officials describe the ASCI software development effort as a hierarchy of development. At the lower level of the hierarchy are software modeling efforts that include (1) modeling the engineering features and the materials used in weapons systems, (2) modeling the physics phenomena associated with weapons systems, and (3) developing computational problem-solving techniques that will allow calculations to take place at increasingly higher processing speeds. At the top of the software development hierarchy are the integrated software applications that will eventually (as larger ASCI computers become available) incorporate all the lower-level modeling efforts and computational techniques into a single system. This integrated software is expected to provide the ability to simulate weapons performance ranging from individual components to full weapons systems, including performance in

³Interim Report to the President. The President's Information Technology Advisory Committee, Aug. 1998.

hostile environments. This software is also expected to provide the capability of predicting the performance of weapons components and full weapons systems in analyses of design, aging effects, and accident scenarios. Program officials also expect that ASCI software will be used to design efficient and environmentally acceptable manufacturing processes.

The report by the ASCI Blue Ribbon Panel noted several concerns about the modeling efforts of the lower-level software. Among the concerns raised was one about the materials science area. The panel said that this area warrants further review because it forms the basis of so much of the work and involves issues of great complexity, some of which are not understood at a fundamental level anywhere in the materials science community. The reviewers also noted that the presentations they heard did not indicate that the integration of experimental data is tightly coupled to software development. The report notes that a robust experimental program that is closely tied to simulations is crucial to assess the adequacy of the scientific input and to test the software.

Another technical challenge in the area of software development is the development and consistent use of software quality assurance. In general, software quality assurance involves reviewing and auditing software products and activities to verify that they comply with the applicable procedures and standards. An April 1998 review of the computing facilities at Los Alamos concluded that software quality assurance has not been addressed sufficiently in the ASCI program.⁴ The report noted that the situation is exacerbated because of the current shortage of expertise in the area of software quality assurance. ASCI program officials have explained that efforts to ensure software quality are part of their new software verification and validation effort and that they have efforts under way at each of the laboratories to address this issue. For example, Livermore has established the Software Technology Center, and a software quality assurance team has been formed using staff from each of the three laboratories. This team recently conducted a survey at each of the laboratories to develop an initial inventory of software quality practices being used at the laboratories. The team is preparing detailed reports on its survey findings that will be provided to each of the laboratories.

Significant Technological Improvements Needed in Some Infrastructure Areas

The major technical risks in the infrastructure area are associated with (1) extracting optimum simulation performance from tens of thousands of processors and (2) moving, storing, and displaying large, complex results

⁴Los Alamos CIC Division External Review Committee Report for the April 1998 Review.

for interpretation by weapons designers. According to program officials, risk mitigation strategies employed in this area include ongoing research at the laboratories and through the Academic Strategic Alliances Program and contracts with third-party partnerships to develop a variety of advanced techniques and technologies. In addition, the ASCI program has sponsored workshops with universities, other government agencies, and industry to engage them in a common approach to meeting these challenges.

Significant technological improvements are needed in several critical infrastructure areas, including visualization and storage technology and the technology that connects computers to other components. For example, visualization is an essential analysis tool for understanding the volumes of data that will be produced by ASCI software. The laboratories have recently unveiled new data visualization centers, but according to ASCI planning documents, the defined user needs exceed industry visualization hardware capabilities by 15 to 60 times. Achieving the needed improvements is also challenging because there is currently only one vendor in this area. Significant improvements will also be needed for data storage technology and connection technology for the 30-trillion operations per second system. Contracts have been signed and work is under way with several vendors to address these issues.

The ASCI Blue Ribbon Panel also cited visualization technology as an area of concern, particularly the level of involvement by weapons designers in the planning of visualization capabilities and facilities. The report also noted that the panel was concerned about the accelerated pace of investment in visualization technology because the basis for visualization needs was unclear.

Integration of Scientific Data

Integrating the data from laboratory experiments conducted outside the ASCI program into ASCI software development efforts has been noted as another important technical challenge. DOE's Green Book notes that to achieve the modeling and simulation goals of the ASCI program, new data will be needed from laboratory experiments to help verify the accuracy of the ASCI software. These experiments are designed to learn more about the physical processes that occur to a weapon under normal and abnormal conditions. The Green Book also notes that the schedule for future experiments and the computational needs of the ASCI program must be closely and carefully coordinated to ensure that the experimental data are useful to the ASCI program. The facility plan for one of the stockpile

stewardship program's experimental facilities illustrates the connection between data from experiments and the ASCI program. The plan for the National Ignition Facility (NIF) explains that much of the program is designed to gather fundamental weapons-relevant data and use these data to enhance and refine nuclear weapons simulations.⁵

Recent reviews have commented on the issue of physics data in the ASCI program. The ASCI Blue Ribbon Panel noted that a robust experimental program, which is closely tied to ASCI simulations, is crucial to assessing the adequacy of the scientific input and to testing the accuracy of the software. The panel reported that the presentations it was given by laboratory officials did not indicate that the experimental data were tightly integrated with software development. Although the panel did not review the experimental program in depth, it recommended that additional funding should be made available to produce the physical data required to support ASCI software efforts. Two June 1998 reports on the software development efforts at Los Alamos expressed concern that the issue of weapons physics had not received the attention it deserved.⁶ One report explained that the review panel wanted to learn more about the role of experiments in validating the accuracy of ASCI software and that certain laboratory staff should have a more prominent role in the selection of experiments conducted at DOE facilities.

Program officials acknowledge the need for closer integration between laboratory experiments and the ASCI program. They explained that they have taken actions such as reorganizing DOE management and creating a formal software validation program that requires data from experiments. In addition, officials explained that the 1998 review of the primary burn code milestone reported on the effective integration of experiments with the ASCI program. According to DOE officials, the reorganization of offices within DOE's Office of Defense Programs, now in progress, will allow ASCI program officials to set both the ASCI program schedule and the schedule for needed laboratory experiments. Program officials also explained that the ASCI "verification and validation" effort, new for fiscal year 1999, would provide the framework for aligning the needs of the ASCI program with the schedule for laboratory experiments. The validation effort includes the use of laboratory experiments to ensure that the simulations are consistent with observed behavior. The June 1998 review of the

⁵Facility Use Plan of the National Ignition Facility, Edition 1, April 1997. NIF is planned as a multiple-beam, high-power laser system with the goal of attaining the ignition of thermonuclear fuel in the laboratory.

⁶"Report of the X-Division Review Committee," May 18-20, 1998, and "Organizational Self-Assessment for the Applied Theoretical and Computational Physics Division," June 26, 1998.

primary burn code milestone found ample evidence of integration among ASCI software development, experiments, and the use of existing information from previous nuclear tests. The review noted that data from experiments are currently being incorporated into ASCI software.

Technical Expertise

Recruiting and retaining qualified personnel is a continuing area of risk, according to ASCI program officials and outside program reviews. DOE and laboratory officials have explained that they have efforts ongoing in many areas to improve the recruiting of staff with the required expertise. These efforts include exemptions from salary ceilings and the ASCI Academic Strategic Alliance Program. The Chiles Commission noted that there is no certainty that DOE will succeed in maintaining future nuclear weapons expertise.⁷

ASCI program officials are concerned about the availability of staff with the necessary expertise. For example, Los Alamos officials noted that several milestones during fiscal year 1998 were delayed because of a shortage of staff with the needed expertise in software development. They explained that such personnel are difficult to recruit and that, once recruited, they need time to develop the necessary weapons-related expertise. Furthermore, these officials said that they have received a waiver from DOE to offer more competitive salaries to recruit qualified staff. Livermore officials explained that personnel with computer science and math skills are in high demand, which makes it difficult to recruit them into the ASCI program. Their risk mitigation strategy includes using the Academic Strategic Alliance Program to attract qualified students, offering competitive salaries, and using the unique research and development aspects of the ASCI program to attract potential candidates.

The Chiles Commission report on maintaining nuclear weapons expertise noted that it was difficult to conclude that DOE will succeed in maintaining future nuclear weapons expertise. Although the report found a great deal that is healthy in the nuclear weapons complex, with many trends moving in the right direction, it also found other matters that are disturbing. These other matters included the aging workforce, the tight market for talent, and the lack of a long-term hiring plan. The report also concluded that steps need to be taken now to ensure that the upcoming generation of designers is recruited and trained while the more experienced designers remain at the laboratory or are available through retiree programs. The

⁷“Report of the Commission on Maintaining United States Nuclear Weapons Expertise.” Report to the Congress and the Secretary of Energy, March 1, 1999.

report explained that more than 60 percent of the nuclear weapons designers at Los Alamos and Livermore are between 50 to 65 years old and that 5 or more years of experience working with experienced designers is required to develop a fully capable, independent designer.

The ASCI Blue Ribbon Panel report concluded that the training of the next generation of technical staff is the single largest problem facing not only the ASCI program but also the entire weapons program. The report commended the ASCI program for its Academic Strategic Alliances Program to attract high-quality, technically trained personnel but also stated that its fellowships and summer internships must be made more appealing and competitive. The report also noted that ASCI computers could be used to attract students into the field of computational science and that, although this is being done in the Academic Strategic Alliance Program, it could be done more broadly. In addition, the report questioned whether the laboratories are able to compete in the market for the best personnel. The report concluded that ties to the universities would be vital to convince the best students to make a career at the laboratories. Program officials acknowledge the problem of recruiting and retaining staff expertise. They cite efforts such as a DOE fellowship program as a means of attracting needed expertise. Under the fellowship program, ASCI will support eight students, a number that may double in the next 2 years. Students must serve a “practicum” at DOE or the laboratories as a condition of support. Program officials hope that this experience will interest the students in working at the laboratories when they complete their education.

Program- and Laboratory-Level Software Milestones From DOE Planning Documents, Fiscal Years 1996-2004

Fiscal year	Milestone description	Program-level documents			Laboratory-level documents		
		Program plan (Sept. 96)	Most critical milestones (July 98)	Green Book ^a (Apr. 98)	Consolidated roadmap ^b (Feb. 98)	Nuclear roadmap ^c (Mar. 99)	Nonnuclear Roadmap ^d (Mar. 99)
96	Microaging				X	X	
97	Microaging	X	X	X			
	High-fidelity safety calculation	X	X	X	X		
	Neutron generator standoff						X
	3-D casting microstructure				X	X	
98	3-D casting microstructure	X	X	X			
	3-D nuclear safety simulation/ Prototype nuclear safety simulation ^e		X			X	
	Neutron generator hostile certification/ Neutron generator radiation hardness mechanical ^e		X				X
	Prototype 3-D physics				X		
	Prototype 3-D hydrodynamics/ radiation-hydrodynamics					X	
	System/composition thermal						X
	B61 penetrator						X
	Macro/micro aging				X		
99	Macro/micro aging	X	X	X			
	Prototype 3-D physics/ Prototype 3-D primary simulations/ 3-D burn code/ 3-D primary burn codes/ 3-D primary burn prototype ^e	X	X	X	X	X	
	3-D forging/welding microstructure	X	X	X	X		
	Crash/fire safety	X		X	X		
	2-D deterministic radiation transport						X
	Abnormal environment thermal assessment						X
	Neutron generator performance code						X

(continued)

**Appendix I
Program- and Laboratory-Level Software
Milestones From DOE Planning Documents,
Fiscal Years 1996-2004**

Fiscal year	Milestone description	Program-level documents			Laboratory-level documents		
		Program plan (Sept. 96)	Most critical milestones (July 98)	Green Book ^a (Apr. 98)	Consolidated roadmap ^b (Feb. 98)	Nuclear roadmap ^c (Mar. 99)	Nonnuclear Roadmap ^d (Mar. 99)
	Full system (Salinas)						X
00	Full physics, full system prototype	X		X			
	Hostile environment electronics certification		X				
	Abnormal environment thermal assessment		X				
	Prototype 3-D secondary simulations/ 3-D secondary burn prototype ^e		X			X	
	Pit casting and manufacturing code				X		
	3-D prototype radiation flow simulation					X	
	Full system microaging simulation					X	
	Parachute						X
	Reentry vehicle aerodynamics						X
	B61 laydown						X
	Electrical circuit simulation capability						X
	Component deterioration model				X		
01	Component deterioration model	X	X	X			
	Initial operating code/ Prototype 3-D coupled simulation/ 3-D secondary burn code ^e	X	X	X	X		
	Stockpile-to-target-sequence certification demonstration				X		
	System composition burn						X
	3-D prototype full system coupled simulation					X	
02	Full physics burn code prototype				X		
	Full system radiation hardness & hostile						X

(continued)

**Appendix I
Program- and Laboratory-Level Software
Milestones From DOE Planning Documents,
Fiscal Years 1996-2004**

Fiscal year	Milestone description	Program-level documents			Laboratory-level documents		
		Program plan (Sept. 96)	Most critical milestones (July 98)	Green Book ^a (Apr. 98)	Consolidated roadmap ^b (Feb. 98)	Nuclear roadmap ^c (Mar. 99)	Nonnuclear Roadmap ^d (Mar. 99)
	Abnormal stockpile-to-target-sequence						X
	Burn code with aging				X		
03	Integrated full physics burn code prototype				X		
	3-D electrical device physics						X
	Abnormal stockpile-to-target-sequence with aging						X
	Normal stockpile-to-target-sequence						X
04	Normal stockpile-to-target-sequence with aging						X
	Complete physics full system prototype/ Initial full system simulation code/ 3-D high fidelity physics full system initial capability ^e		X		X	X	

^aThe full title of this document is Stockpile Stewardship Plan - 2nd Annual Update.

^bThe full title of this document is Consolidated Applications Roadmap.

^cThe full title of this document is Consolidated Nuclear Component Applications Roadmap.

^dThe full title of this document is Non-Nuclear Codes Roadmap.

^eThis represents a single milestone that is titled differently in multiple sources.

Sources: ASCI Program Plan (1996), Stockpile Stewardship Plan (1998), and other information provided by the Department of Energy, and the Los Alamos, Livermore, and Sandia National Laboratories.

Comments From the Department of Energy

Note: GAO comments supplementing those in the report text appear at the end of this appendix.



Department of Energy

Washington, DC 20585

June 7, 1999

Ms. Susan D. Kladiva
Associate Director, Energy, Resources,
and Science Issues
U.S. General Accounting Office
Washington, D.C. 20548

Dear Ms. Kladiva:

We have reviewed the draft of your proposed report entitled NUCLEAR WEAPONS: DOE Needs to Improve Oversight of the \$5 Billion Strategic Computing Initiative (GAO/RCED-99-195, Code 141213), which you transmitted to us on May 26, 1999.

We are pleased to read that the General Accounting Office (GAO) affirms the Department of Energy's (DOE) approach to the scientific challenge undertaken by the Accelerated Strategic Computing Initiative (ASCI). In this draft report, the GAO expresses support for the urgency of ASCI's principal goal to establish a simulation-based approach to maintain the Nation's nuclear deterrent. The draft report also recognizes the critical importance of the hardware and software technologies that ASCI is pursuing to achieve its goal and the significant technical challenges presented by this goal. The draft GAO report further confirms the importance of ASCI's aggressive platform development and software development plans. The GAO observes that the use of smaller and less capable computers is no longer feasible in light of the more complex and detailed simulations necessitated by the requirements of the Stockpile Stewardship Program. The GAO also acknowledges the technical challenge that the Department faces in developing the highly technical and complex applications software needed to predict weapons performance accurately.

This report states that, "ASCI computers and software developed to date are already important tools for addressing DOE's high-priority stockpile needs." The Department appreciates the recognition by the GAO of ASCI's progress in meeting the difficult technical challenges it faces in this complex undertaking.

After careful review, the Department is providing the enclosed response to the draft report. Although the Department agrees with many of your findings, and has accepted some of the recommendations, there are areas and recommendations which we disagree with and non-concur. It should be noted that GAO officials have been informed that this draft report contains procurement sensitive information, and this information stated in Chapter 3, Unreliable Cost Estimates Also Contribute to Cost Increases, should be removed.

Appendix II
Comments From the Department of Energy

Thank you for the opportunity to provide comments on the draft report. If you require additional assistance or have questions about the comments, please refer them to Dr. Paul Messina of my staff at 202/586-1101.

Sincerely,



Victor H. Reis
Assistant Secretary
for Defense Programs

cc:
Mike Telson, CFO
Betty Smedley, CR-1
Gilbert Weigand, DP-10
Jim Landers, DP-40
Linda Thomas, DP-44

**DOE Comments on the Draft General Accounting Office Report
“NUCLEAR WEAPONS DOE Needs to Improve Oversight of the \$5 Billion
Strategic Computing Initiative”**

General Comments

We are pleased to read that GAO affirms the Department of Energy's (DOE's) approach to the scientific challenge undertaken by the Accelerated Strategic Computing Initiative (ASCI). In this draft report, the GAO expresses support for the urgency of ASCI's principal goal to establish a simulation-based approach to maintaining the nation's nuclear deterrent. The draft report also recognizes the critical importance of the hardware and software technologies that ASCI is pursuing to achieve its goal and the significant technical challenges presented by this goal. The draft GAO report further confirms the importance of ASCI's aggressive platform development and software development plans. The GAO observes that the use of smaller and less capable computers is no longer feasible in light of the more complex and detailed simulations necessitated by the requirements of the Stockpile Stewardship Program. The GAO also acknowledges the technical challenge that the Department faces in developing the highly technical and complex applications software needed to predict weapons performance accurately.

This report states that, “ASCI computers and software developed to date are already important tools for addressing DOE's high-priority stockpile needs.” The Department appreciates GAO's recognition of ASCI's progress in meeting the difficult technical challenges it faces in this complex undertaking.

After careful review, the Department provides the following response to the recommendations presented in the report.

GAO RECOMMENDATION

DOE has chosen not to designate the strategic computing program as a strategic system. Given the strategic computing program's estimated cost of over \$5 billion, the lack of a comprehensive planning, tracking, and reporting system, and the importance of the program to maintaining the stockpile of nuclear weapons, it is important that DOE improve its oversight and management of this program. Therefore, we recommend that the Secretary of Energy require the establishment of a comprehensive planning, progress tracking, and reporting system and that the Secretary designate the program as a strategic system warranting oversight at the highest departmental level.

CONCUR: The Department concurs with the first part of this recommendation. DOE has established a comprehensive planning, progress tracking, and reporting system. In the early stages of the program, comprehensive planning and progress tracking were implemented through issuance of an ASCI Program Plan and Implementation Plans. An updated version of the ASCI Program Plan will be issued very soon which contains detailed specifications for all of the critical programmatic milestones, and the FY2000 ASCI Implementation Plans will be completed by September 30, 1999. A detailed Work

Now on p. 6.
See comment 1.

Breakdown Structure has been developed for these plans, ensuring that they will contain descriptions of all program elements and complete lists of all programmatic milestones. A new progress reporting mechanism has already been established. The weapons laboratories have submitted their first quarterly progress reports. These reports address every programmatic milestone due in FY1999 and permit DOE managers to track the progress of their projects.

NON-CONCUR: The Department does not concur with the recommendation that the Secretary designate ASCI as a strategic system. In accordance with DOE Order 430.1A, a strategic system is primarily intended for life cycle management for physical assets. ASCI is a program not a physical asset as defined in the DOE order. To designate ASCI as a strategic system is not appropriate for the management of this program. We believe that the planning, progress tracking, and reporting systems that have been implemented are appropriate. To implement the recommendation would create a duplicative non value-added process. ASCI is recognized as a major national security mandated effort and currently undergoes extensive internal and external reviews as part of the overall management of the Stockpile Stewardship Program. DOE has a comprehensive secretarial review process that meets the intent of the Clinger-Cohen Act of 1996, and this program receives the high-level oversight it requires. We are confident that the DOE has an effective, efficient, and appropriate management process with the flexibility required by a leading edge developmental program.

GAO RECOMMENDATION

Furthermore, given the substantial increases in the cost estimates for the strategic computing initiative to date, DOE's weaknesses in estimating costs for the unprecedented scale of development efforts, and the lack of a project cost tracking process, we also recommend that the Secretary require the strategic computing initiative to adopt systematic project cost tracking procedures that will allow DOE managers to determine if specific projects are within budget.

NON-CONCUR: Almost all of the referenced increase was due to the considerable expansion in scope of the ASCI program in order to compensate for the ban on nuclear testing imposed by the Comprehensive Test Ban Treaty. These additional program elements were reviewed, validated and the President and the Congress agreed that they should be added to the overall program. Thus, the costs have increased only minimally against the original size, scope and timeframe of the program. The program has been refined and rebaselined, and the costs for this expanded and enhanced program have been identified and presented in the appropriate budget. The ASCI requirements of this change were considerable, due to the significantly increased dependence on the physical simulations, resulting in additional need for computing resources. In July 1997, the program completed a zero-based budget exercise to estimate program costs. Since then the program's cost estimates have changed very little.

The Department's cost tracking procedures for ASCI projects are consistent with and appropriate for the fundamentally developmental nature of this important effort. Funding and costs are tracked and reported at the strategy level using Budget and Reporting

(B&R) classification codes and the Financial Information System (FIS), the official Departmental financial system. The tracking systems are extended down to the individual projects with other funding and cost monitoring tools that gather more detailed information. For example, in March we conducted an analysis of three ASCI/SC strategies that identified commitments and cost status for specific lab procurements at the project level. In addition, the Department is currently instituting enhanced processes and laboratory coordination as an integral part of its more comprehensive planning, tracking and reporting system. These tools will augment the formal cost tracking and reporting systems in an efficient and effective manner. Because of the developmental nature of the ASCI program, some budgetary flexibility is necessary to capitalize on or adapt to changes within the high-performance computing industry.

Specific Comments

Below are the Department's comments to specific GAO findings.

GAO Finding
Executive Summary, Technical Challenges Are Present In All Aspects of the Strategic Computing Initiative, Page 7

"Nonetheless, as noted in one review of the program, there is no certainty that DOE's efforts will succeed."

DOE Response
GAO's comment that there is no certainty that we will succeed can be applied to any program, particularly a leading edge developmental program. ASCI has a strong track record of success since its inception in 1996, and there is every indication that ASCI will succeed in the end. ASCI is ambitious and every major programmatic milestone has been met in the advancement of computers, modeling, and science. Every technical review has affirmed that ASCI's goals are achievable. ASCI advances have already made significant contributions to weapons stewardship and to science and engineering in general.

GAO Finding
Chapter 2, Comprehensive Planning is Needed, Page 17

"The ASCI program's long- and short-term planning efforts thus far have been inconsistent and incomplete."

DOE Response
We have had implementation plans since the inception of this program in 1996, and these plans have been improved consistently each year since. A technical review of ASCI by the Blue Ribbon Panel dated March 2, 1999, came to a very different conclusion than did the GAO. This panel consisted of distinguished experts from academia, government, the nuclear weapons laboratories, and the private sector, and it was chaired by Dr. Venkatesh Narayanamurti, Dean of Engineering and Applied Science at Harvard University. They stated that "We are impressed with the good job that the ASCI program management structure is doing managing a large, complex, time-critical program. We encourage you

Now on p. 6.
See comment 2.

Now on p. 19.
See comment 3.

Appendix II
Comments From the Department of Energy

to continue developing resource plans through tri-lab analysis of technical and programmatic requirements. The plans are dynamically managed to identify any missing program elements early, and to reallocate resources to maintain program schedules.”

GAO Finding

Chapter 2, Strategic and Annual Plans Are Outdated..., Page 17

“Although hardware and infrastructure milestones have not changed much, program officials have revised software milestones numerous times but have not published an updated strategic plan.”

DOE Response

It is true that software milestones are revised on occasion. However, most of these revisions do not affect the strategic direction of the program. Software project plans by necessity are living documents, reviewed and revised frequently as research results come to light, as tasks are completed, or as tasks are no longer needed for technical reasons. Furthermore, high-level software milestones are created from the program strategic plans and other planning activities. These milestones are created specifically to meet the strategic needs of the Stockpile Stewardship Program and are not changed on a continual basis. Lower-level milestones and tasks, however, are changed as needed and are tracked by each project to ensure accurate timetables for meeting high-level milestones. This is all part of the management process of a technical program.

GAO Finding

Chapter 2, Strategic and Annual Plans Are Outdated..., Page 18

“For example, the “abnormal environment thermal assessment” milestone is listed as occurring in fiscal year 1999 in one planning document and in fiscal year 2000 in another. (App. I lists ASCI program milestones as identified in various planning documents),”

DOE Response

In both milestone charts, the milestone is depicted on the line that separates FY99 from FY00. In one chart, the milestone is centered on the line, and on the other chart, the milestone mark has been moved slightly into FY00. Our written documentation on this milestone clearly states that it is due at the end of the 4th quarter of FY99. Furthermore, the GAO requested clarification of this milestone in an email request of May 1, 1999. We responded that the milestone was a fourth-quarter, FY99 milestone. There should have been no confusion about this milestone.

GAO Finding

Chapter 2, Strategic and Annual Plans Are Outdated..., Page 19

“As with other ASCI program documents, these sources report milestones inconsistently and laboratory-level milestones may or may not match program-level milestones.”

Now on p. 19.
See comment 4.

Now on p. 20.
See comment 5.

Now on p. 21.
See comment 6.

Appendix II
Comments From the Department of Energy

DOE Response

Milestones that appear in “The Most Critical Milestones” chart are very high-level program milestones that represent a consolidation of lower level milestones appearing in the Implementation Plans. Milestones that appear in “The Most Critical Milestones” chart are not identical to intermediate-level milestones on the three charts “Consolidated Roadmap,” “Nuclear Roadmap” and “Non-Nuclear Roadmap,” which are more detailed, but still consolidations of milestones in the Implementation Plans. These charts were created to provide strategic program information at different levels of detail and should not be expected to show identical milestones.

Furthermore, these charts were not meant to provide details about milestones. These are visual supplements to written documentation that provides detailed information about the milestones. Transcription and printing errors in the electronic version of the Consolidated Roadmap chart (which provided the basis for the Nuclear Roadmap) caused the labels on the abscissa, which assigned dates to the milestones, to shift so that several incorrect versions of the chart were printed and made available to the GAO. In each of those versions, the charts are identical except for the shifted abscissa labels. This simple problem explains virtually all the contradictory information about milestones detailed in Appendix I. The GAO auditors should have recognized the error and requested a correction of the chart. In no case was the project unclear about the details of the milestones on these charts in our verbal communications or our detailed documentation.

GAO Finding

Chapter 2, The Linkages Between Annual and Long-Term Milestones..., page 20

“With the exception of information provided in response to our request, these linkages have never been documented.”

DOE Response

While we agree that linkages are not always clear to the GAO, it is not correct to state that linkages have never been documented nor is it correct to infer that technical staff do not understand the linkages. Again, the Blue Ribbon Panel came to a completely different conclusion based on their technical review of the program. The milestones, both long- and short-term, are used as planning and management tools at all of the laboratories. In fact, annual and many long-term milestones are explicitly present in individual ASCI code project-planning documents. It is these milestones that drive the technical development of the application projects themselves. For example, the CY99 burn code milestone appears in the primary code projects’ planning documents and specific tasks and low-level milestones are formally linked to this milestone. This milestone is further linked to longer-term milestones as a building block to reach the longer-term goals. These linkages are clear to management and technical staff alike.

Furthermore, the GAO claims in this report that the “principal investigators” (PI) meetings are not helpful in ascertaining the accomplishments made on individual projects as related to the annual or long-term milestones. This is a misunderstanding. The presentation templates for these meetings are generated precisely for the purpose of

Now on p. 22.
See comment 7.

uniform tracking of project progress against such milestones, and in fact, must be filled in with milestone-relevant information. These meetings are lengthy, highly technical, and, like the planning documents discussed above, are not intended to be tutorials for the lay person. It is likely that non-technical individuals in the weapons program would find information presented at these meetings difficult to follow.

GAO Finding

Chapter 2, The Linkages Between Annual and Long-Term Milestones..., Page 20

“At Sandia, for example, laboratory officials identified ASCI software activities that are needed to meet stockpile requirements but that are not tied directly to program-level milestones.”

DOE Response

There seems to be considerable confusion regarding the relationship of program-level milestones and the project-level milestones throughout the draft report. Sandia identified detailed ASCI software project activities that are needed to meet stockpile requirements, but these are not necessarily program-level milestones. Program-level milestones represent a consolidation of lower level milestones and activities. Milestones are reported at different levels depending on whether the program or a particular project is being discussed.

GAO Finding

Chapter 2, Plans Should Clearly Identify Research Strategies..., Page 21.

“In the area of infrastructure development, for example, one overall goal is to develop ways for scientists to examine massive amounts of weapons simulation data. To achieve this goal, the development effort has seven simultaneous lines of effort, including developing “common data formats” and “distributed file systems.” A program official explained that all of these lines of effort are needed but that it is not clear at this time which efforts will be critical to meeting the overall goal. It is also not known when a choice among these options needs to be made.”

DOE Response

The examples cited above are neither interchangeable nor mutually exclusive. The development of common data formats is recognition that, in order to facilitate the examination of data from different calculations at different sites, it would be much easier to exchange that data if a common format for sharing the data were developed. Such a format has existed for some time, but an update is required for ASCI-scale problems. The development of “distributed file systems” technology also relates to data sharing and manipulation at remote sites, but from a security and hardware-level-performance perspective. These are not redundant, risk-mitigation paths, but completely separate program elements to address separate technical needs that exist within the realm of sharing large data sets. Naturally, risk-mitigation paths do exist, but these are not among them.

Now on p. 22.
See comment 8.

Now on p. 23.
See comment 9.

The general conclusion drawn, that it is not clear when a choice among these options needs to be made, is accurate. Some of the technology being developed is new and, as a result, risky. Risk-mitigation paths exist to alleviate this risk. Many examples of risk-mitigation paths in which one path was abandoned when another path proved fruitful exist within the ASCI program. However, these are typically not on a set timetable, because it is difficult to predict when one research area will bear fruit over another. This is the inherent nature of scientific research and development. Again, quoting from the Blue Ribbon Panel report, "The program correctly manages risk through pursuit of more than a single approach to more difficult tasks, such as those identified in the barrier analysis."

GAO Finding

Chapter 2, Plans Should Define Performance Criteria for Milestones, Page 23

"Program officials acknowledge that they need to perform similar reviews for other milestones. At this juncture, however, most program-level milestones remain undefined."

DOE Response

Similar reviews are being performed for other milestones. During the GAO investigation, an external review of the non-nuclear mechanics codes was completed. This review entitled "Codes for the Complex" was conducted in April 1999 and a report was issued in May 1999. The panel concluded that the code development efforts were focused on the right technological areas, the development targets were realistic, and that the quality of the work was excellent and appropriate progress was being made.

GAO Finding

Chapter 2, Systematic Progress Tracking and Reporting Needed, Page 23

"The ASCI Program Plan described the semiannual Principal Investigators meetings as the primary forum for reporting program progress...While these meetings are a forum for ASCI researchers to exchange ideas, there are few reporting requirements, and reports and program accomplishments have generally not been related to established milestones."

DOE Response

While it is true that the meetings themselves focus on individual projects and are an excellent forum for ASCI researchers to exchange ideas, it is not true that there are few reporting requirements or that accomplishments are not related to established milestones. These meetings are highly technical, detailed, and lengthy. Each project must strictly adhere to a reporting format, complete with timelines and completions of individual tasks as related directly to established milestones. That is the fundamental purpose of these meetings. As stated above, these meetings are not designed for consumption by the general public and, in order to cover the large number of projects in ASCI, presentations are generally constrained to present the information in highly compressed technical form.

Now on pp. 24-25
See comment 10.

Now on p. 25.
See comment 11.

Now on p. 25.
See comment 12.

GAO Finding

Chapter 2, **Systematic Progress Tracking and Reporting Needed**, Page 23

“For example, accomplishments reported for fiscal year 1998 for one key software project did not reveal that the project was actually 6 months behind schedule, which affected the schedule of other related projects.”

DOE Response

This statement is a misinterpretation of the information provided to GAO. Sandia stated that relative to the projection made at the beginning of the program four years ago, the software project in question was approximately six months behind schedule. However, the program schedule was rebaselined in the beginning of FY98 and since that time the project has remained on schedule. Therefore, the report for FY98 was accurate and the project was not six months late. Furthermore, the schedule of this project is not adversely impacting the schedules of other code projects. Schedule risks are typically anticipated when working on a science research and developmental program and tasks within projects are adjusted when new and relevant information is discovered.

GAO Finding

Chapter 2, **Systematic Progress Tracking and Reporting Needed**, Page 24

“To varying degrees, the laboratories used the annual implementation plans to report accomplishments during fiscal years 1996 through 1998, although this practice was discontinued in fiscal year 1999.”

DOE Response

This is a misunderstanding of the information provided to GAO. The implementation plan (IP) process has been strengthened each year, becoming a budget quality document in FY99. The Laboratories now provide quarterly (rather than yearly) IP progress reports to DOE headquarters. These quarterly progress reports are now the preferred vehicle for reporting of accomplishments.

Now on p. 26.
See comment 13.

GAO Finding

Chapter 2, Hardware Performance, Page 25

“However, Los Alamos is still working to address problems affecting the reliability and stability of its system that, as a result of these problems, did not pass [sic] all its contractually required acceptance tests on time.”

DOE Response

The computer system did not pass all of its originally contractually required acceptance tests on time. However, LANL negotiated a stability plan with one of our industry partners, Silicon Graphics Inc. (SGI), and tied this plan to the remediated contract, and is in the process of modifying the milestone payment schedule to reflect this plan. Current data indicates good progress toward full system stability.

Now on p. 27.
See comment 14.

Appendix II
Comments From the Department of Energy

It is again appropriate to quote from the Blue Ribbon Panel Report that praised the ASCI hardware program: "The hardware program is very impressive, and is clearly a world class effort. It has pushed the computer industry to go significantly beyond where it would otherwise be. At the same time, it has not gone so fast that the industry has been unable to produce integrated systems or the software to go with them. As a result, ASCI has had the important secondary effects of forwarding the high-performance computer industry, and the field of computer science... We are convinced that the program management understands the problem, and we urge that the schedule be reconsidered regularly."

We anticipated stability issues in standing up the largest computational platform in the world. We planned all along to work closely with our industrial partner to address issues as they arose. The quote from the Blue Ribbon Panel is recognition of this effort and a reminder of the teamwork required to execute a difficult technical program.

GAO Finding
Chapter 2, Hardware Performance, Page 25

"Progress in hardware development has been reported in annual implementation plans and at semiannual principal investigator meetings and in press releases to the media... These releases may be misleading in some cases because they are reporting on only selected aspects of performance. For example, DOE reported that the Los Alamos system established a world record for sustained speed. However, it was not reported that the test used was substantially easier than the test specified in the contract or that the system has not yet met other acceptance test criteria (such as, mean time between failures). The sustained speed of this machine would be one-half the speed reported in press releases had the test specified in the contract been used."

DOE Response

This statement is not true on several counts.

First, the one trillion operations per second (TeraOPS) sustained performance using the sPPM application was required as part of the original contract. However, this contract was remediated some time ago when SGI canceled a processor necessary to achieve this performance and it is this remediated contract that Los Alamos and SGI are currently working from. The remediated contract, based on SGI's strategic plan and continuing technical work on the Los Alamos system, required a "best effort" on the sPPM benchmark, which has been done.

Second, it is not correct to state that the sustained speed of this machine would be one-half the speed reported in press releases had the test specified in the original contract been used. In fact, the performance of this benchmark greatly varies as a function of the size of problem run.

Third, the successful execution of Linpack and the record setting performance are all critical elements of the stability and performance assessment of the Los Alamos system.

Now on p. 27.
See comment 15.

It is misinterpretation to state that Linpack is substantially easier than any other “planned” tests. Linpack is in fact a relevant test of the computational hardware, as it closely resembles components of the actual ASCI applications. It is also an industry standard used to assess computational performance. The Linpack run was, in fact, a critical success leading to system stability and is an important part of the joint SGI/Los Alamos stability plan in the remediated contract.

Finally, it is important to emphasize that stability problems were expected and anticipated. Standing up the largest computational platform in the world is inherently difficult. Such systems are initially inherently unstable. The stability of the Los Alamos system has improved significantly over time. Working together, the ASCI applications, our industrial partner (SGI), and local computing support at Los Alamos have greatly improved the stability of the system.

GAO Finding

Chapter 2, Hardware Performance, Page 25.

Of the 100,000-fold increase in computing performance needed by ASCI, DOE expected a 10-fold increase from taking advantage of parallel-processing techniques.

DOE Response

This statement is puzzling to us. The DOE is buying computers with approximately 10,000 processors. With good parallel scaling efficiency, we plan to achieve a substantial fraction of a 10,000-fold increase in performance from parallel-processing techniques. Algorithmic advances are expected to contribute part of the desired performance increase. The remaining factor of roughly 10 would come from increased individual processor speeds attained during the course of the program.

GAO Finding

Chapter 2, Software Development, Page 26.

“Documentation of ASCI’s software performance is limited. With the exception of the “Prototype 3D Primary Simulations” milestone review discussed previously, little information exists that evaluates software development efforts.”

DOE Response

This is a misunderstanding of the information provided. All of the ASCI projects, particularly the application projects, have numerous peer-reviewed articles in scientific journals and conference proceedings. One project alone has over 30 publications during just FY99. In addition, most of the projects participate in peer-review activities with other software development sites and individuals with expertise in the field. Los Alamos has had many industry recognized and respected software technical advisors and visitors from around the world talking with the software project leaders, and all of the ASCI projects have participated in an independent, formal quality review based on the Software Engineering Institute’s (SEI) Capability Maturity Model (CMM).

Now on p. 27.
See comment 16.

Now on p. 27.
See comment 17.

Appendix II
Comments From the Department of Energy

Quite contrary to this observation, there is ample information of software development, performance, and assessment on the ASCI projects available. Again, quoting from the Blue Ribbon Panel report dated March 2, 1999: "The approach to software development is very sensible. The abstraction-based tools being developed are effective means for achieving the needed code efficiency and portability. The design of the software environment is as important as code creation, and we are convinced that the effort being put into it will pay long term dividends."

GAO Finding
Chapter 2, Software Development, Page 26

"Our analysis of planning documents, accomplishment reports, and other related material on software activities at Sandia, showed that the ASCI program's current planning and tracking system does not always provide an accurate picture of progress."

DOE Response

This ignores the fact that the program was rebaselined in the beginning of FY98. The current planning and tracking system uses this baseline and provides an accurate picture of progress.

GAO Finding
Chapter 2, Software Development, Page 26

"Sandia officials explained that, although the reported information was accurate, milestones had not been established for all of the important tasks associated with that project."

DOE Response

The statement is a misinterpretation of the information provided. Since the program was rebaselined in the beginning of FY98, the milestones for all important tasks on this project have been established. The FY99 Implementation Plan clearly shows these milestones.

GAO Finding
Chapter 2, Software Development, Page 26

"Since the purpose of this project is to provide a common framework for all Sandia ASCI software models, delays in this project are hindering the performance of other software projects that require integration into the common framework."

DOE Response

This statement is a misunderstanding of the information provided to GAO. The performance of the other software projects has not been hindered due to any schedule changes. Management has anticipated schedule risks and reordered the sequence of tasks as necessary to accommodate any changes.

Now on p. 28.
See comment 18.

Now on p. 28.
See comment 12.

Now on p. 28.
See comment 12.

Appendix II
Comments From the Department of Energy

Now on p. 28.
See comment 12.

GAO Finding
Chapter 2, Software Development, Page 26

“Sandia also did not report these delays in its list of accomplishments for fiscal year 1998 that it submitted to DOE.”

DOE Response
As discussed above, the schedule was rebaselined in the beginning of FY98. Hence, the project in question was not late during the FY98 reporting period.

Now on p. 28.
See comment 19.

GAO Finding
Chapter 2, Software Development, Page 26

“Sandia officials told us that, as of March 1999, 50 percent of their software development projects were experiencing delays due to funding or program changes or are behind schedule for other reasons.”

DOE Response
This is a misunderstanding of the information provided. Sandia stated that relative to the projection made at the beginning of the program four years ago, the Sierra framework project was approximately six months behind schedule. Since many of the software projects will use the Sierra framework, the GAO may have inferred that these projects are behind schedule. This is not the case. The schedule of the Sierra project is not adversely impacting the schedules of other software projects that are on track.

Now on p. 28.
See comment 20.

GAO Finding
Chapter 2, DOE May Need to manage ASCI as a Strategic System, Page 27

“DOE may not be appropriately managing the ASCI program by not designating it as a strategic system.”

DOE Response
To implement the recommendation would create a duplicative non value-added process. ASCI is recognized as a major national security mandated effort and currently undergoes extensive internal and external reviews as part of the overall management of the Stockpile Stewardship Program. Processes that meet the intent of the GAO recommendations are currently in place. We are confident that the DOE has an effective, efficient, and appropriate management process with the flexibility required by a leading edge developmental program.

As part of the budget development process, all proposed investments in scientific mission specific resources are evaluated based on criteria such as life cycle costs, alternative technologies, procurement alternatives, technical risks, impact on scientific research mission, competing mission priorities, etc. Evaluation of ongoing scientific computing projects is integrated with ongoing performance measurement of the research portfolio. Program Secretarial Officers develop and present appropriate justifications and rationale for specific programs at the Secretarial level. After the internal review process, external

Appendix II
Comments From the Department of Energy

reviews by the Office of Management and Budget (OMB), Congressional Committees, and other oversight organizations are done.

In addition to budget reviews, numerous technical internal and external peer reviews of the ASCI program are done. Departmental and tri-lab (Los Alamos, Livermore, and Sandia labs) expertise is supplemented by outside experts in validating procurement requirements and specifications. A JASON review was conducted in 1995. JASON is a DOD sponsored group administered by MITRE Corp. that does independent reviews and evaluations of national security systems, much like the Defense Sciences Board. Many of the recommendations resulting from the JASON review have been implemented in the ASCI program. A DOD Policy and Program Evaluation Office review has been conducted with OMB participation, and a Blue Ribbon Panel of industry and academic experts in high-end computing has endorsed ASCI's platform strategy as a "world class effort". In addition, DOE has a comprehensive secretarial review process that meets the intent of the Clinger-Cohen Act of 1996, and this program receives the high-level oversight it requires.

GAO Finding
Chapter 3, **Cost Estimates Increased Substantially**, Page 29

"DOE's actual fiscal year 2000 budget request for ASCI, which totals \$692 million, is more than double the original fiscal year 2000 estimate made in 1995."

DOE Response

Almost all of the referenced increase was due to the considerable expansion in scope of the ASCI program in order to compensate for the ban on nuclear testing imposed by the Comprehensive Test Ban Treaty. These additional program elements were reviewed, validated and the President and the Congress agreed that they should be added to the overall program. Thus, the costs have increased only minimally against the original size, scope and timeframe of the program. The program has been refined and rebaselined, and the costs for this expanded and enhanced program have been identified and presented in the appropriate budget. The ASCI requirements of this change were considerable, due to the significantly increased dependence on the physical simulations, resulting in additional need for computing resources. In July 1997, the program completed a zero-based budget exercise to estimate program costs. Since then the program's cost estimates have changed very little.

GAO Finding
Chapter 3, **Better Oversight of Costs Needed**, Page 33

"For example, as noted in chapter 2, one laboratory estimates that 50 percent of its software development projects are experiencing delays or are behind schedule."

Now on p. 31.
See comment 14.

Now on p. 35.
See comment 19.

Now on p. 35.
See comment 21.

DOE Response

This is a misunderstanding of the information provided to GAO. To our knowledge, this statement was never made.

GAO Finding

Chapter 3, **Conclusions**, Page 33

“ASCI costs have increased substantially because of changes in program requirements and weaknesses in DOE cost estimation. Because its tracking of costs is limited, DOE cannot determine whether specific projects are under or over budget. Historically, DOE has had difficulty managing the cost of large programs.”

DOE Response

The appropriateness of tracking costs on ASCI projects is a question that deserves careful attention. At its core, the ASCI program is a research and development program and contains elements of both basic and applied research. As such, there must be adequate flexibility and contingency in the program to properly manage the technical risks associated with a program of this type.

Costs are tracked, but not at the individual project level. This management approach is intentional and essential if the program goals are to be met. The approach maintains accountability at the major program element level, while allowing flexibility at the project level. Project level funding decisions and the establishment of project level priorities have been delegated to laboratory management with concurrence by DOE.

This management approach used on the ASCI program has a firm basis and has been recommended by a recent National Academy of Sciences report (“Evaluating Federal Research Programs”). In this report, the NAS states that the methods for measuring progress depend on the nature of the research. There is a distinct difference on how progress should be measured on applied research programs versus basic research programs. Applied research programs should establish specific milestones and outcomes and progress should be measured on an annual basis. For example in the ASCI program, the procurement of supercomputers is managed as an applied research program, where specific deliverables and performance objectives are prescribed in the contract with the vendor. Success is determined by the degree to which these are attained.

For basic research programs, NAS states “it is necessary to evaluate the performance of basic research programs by using measures not of practical outcomes, but of performance, such as generation of new knowledge, the quality of research, the attainment of leadership in the field, and the development of human resources.” Since a substantial part of the ASCI program is research, these latter performance measures are most appropriate. It is for this reason that the Blue Ribbon Panel was convened and these exact questions were used as the basis for their review. While the panel identified some areas for improvement, it stated that “We are impressed that the ASCI program management structure is doing a good job of managing a large, complex, time-critical

program ...The program correctly manages risk through pursuit of more than a single approach to more difficult tasks.”

GAO Finding

Chapter 4, Computer Speed Must Increase Dramatically by 2004, Page 36

“These officials explained that getting computer systems of this size to operate as a fully integrated system has never before been achieved and is one of the most difficult challenges facing the program.”

DOE Response

The GAO fails to mention a major accomplishment of the program: the development of the ASCI Red machine and our unprecedented success in developing a highly reliable and well-utilized production computer. Calculations on ASCI Red machine routinely require the machine to “operate as a fully integrated system” and have allowed researchers to perform important calculations never before possible. Furthermore, ASCI Red was developed on schedule and within budget. The program is already benefiting from early migration of the Blue machines toward stable production platforms in much the same manner and with similar timing as ASCI Red.

GAO Finding

Chapter 4, Computer Speed Must Increase Dramatically by 2004, Page 36.

“An April 1998 review of the computing division at Los Alamos by an external committee recognized this issue by pointing out that users generally had access to only small parts of the computer and rarely had access to the full system. The Committee’s report explained that operating the computer as a fully integrated system was important because of the ASCI computer needs are based on running simulations that require the full capability of the computer. The challenge continues today with the 3-trillion operations per second computer at Los Alamos, which has experienced many failures when trying to run as a fully integrated system.”

DOE Response

Things have changed significantly since 1998. In 1998, the system had only 1024 processors; it now has 6144 and it is feasible to allocate much larger numbers of processors to individual jobs. Again, the stability of the system and the process by which the system stability is being enhanced come as no surprise. See previous discussion of these issues.

Stability has progressed to the point where most of the processor time on ASCI Blue Mountain is devoted to jobs in range of 1000 to 2000 processors. Furthermore, a special queue that runs jobs on the entire 6144 processor system runs each week, and full system tests are also run weekly.

GAO Finding

Now on p. 38.
See comment 22.

Now on p. 38.
See comment 23.

Now on p. 40.
See comment 24.

Appendix II
Comments From the Department of Energy

Chapter 4, Software Development Is Critical to Program Success, Page 38

“The ASCI Blue Ribbon Panel report noted several concerns about the lower-level software modeling efforts. Among the concerns raised was that the materials science area warrants further review because it forms the basis of so much of the work, and involves issues of great complexity, some of which are not understood at a fundamental level anywhere in the materials science community. The reviewers also noted that they did not find evidence to indicate that the data from experiments are tightly integrated with software development. The report notes that a robust experimental program that is closely tied to simulations is crucial to assess the adequacy of the scientific input and to test the software.”

DOE Response

The GAO statement that *“The reviewers also noted that they did not find evidence to indicate that the data from experiments are tightly integrated with software development”* is an incorrect citation. The Blue Ribbon Panel stated that *“The presentations we heard did not indicate that the integration of experimental data is tightly coupled to code development. This may only be our perception, but if it is true, it needs to be addressed.”* While the words are only slightly different, the meanings are quite different.

In spite of these differences, we agree that the coupling of the experimental program to code development is important and should continue to be an integral part of the ASCI program. In the initial stages of the ASCI code development, the emphasis has been on hydrodynamic and radiation transport modeling and methods development. Current code development emphasis on weapons materials and physics models is increasing. The increase in Verification and Validation (V&V) efforts planned for FY99 and FY00 will more closely couple the experimental programs at the laboratories and the ASCI codes. The addition of new hydrotest and pulsed power facilities like DARHT, Atlas and Z-machine, the NIF, Nova, and Maven facilities, and the continuing upgrades to LANSCE will increase the quality, quantity, and timeliness of experimental data delivered to the ASCI program. In fact, we anticipate that the increased V&V activities for existing and new physics models will influence the design and implementation of specific experiments at these facilities.

The organizational changes in DP (DP-10 and DP-50 merger) will facilitate the integration of the experimental and the ASCI programs. The new organization is currently working with the laboratories to undergo a strategic realignment of programs in science and simulation into “campaigns” that support critical program goals. The emphasis of the new organization will be on strengthening the integration of the science and simulation programs.

GAO Finding

Chapter 4, Software Development Is Critical to Program Success, Page 39

Now on p. 40.
See comment 25.

Appendix II
Comments From the Department of Energy

“This team recently circulated a questionnaire to staff at each of the laboratories to develop an initial inventory of software quality practices being used at the laboratories.”

DOE Response

The baseline survey of software engineering practices at the three laboratories performed under the Verification and Validation program was much more than the circulation of a questionnaire. This survey was performed by a team of software quality assurance experts that are trained in the Capability Maturity Model of the Software Engineering Institute. This team spent several weeks at each laboratory and interviewed members of all ASCI code project teams. They also examined project documents and interviewed project managers. As part of this process, each ASCI code project filled a survey. The review team wrote a plan for this survey and is preparing detailed reports on their findings. These reports will be provided to the code teams at all three laboratories.

GAO Finding

Chapter 4, Significant Technological Improvements Needed..., Page 39

“The ASCI Blue Ribbon Panel also cited visualization technology as an area of concern, particularly the level of involvement by weapons designers in the planning of visualization capabilities and facilities. The report also noted that the panel was concerned about the accelerated pace of investment in visualization technology because the basis for visualization needs was unclear.”

DOE Response

Some misunderstanding may have occurred related to the comments made by the designers at the Blue Ribbon Panel review that the GAO references. In fact, in the three workshops on data manipulation and visualization of large-scale scientific data sets that were held in 1998, both primary and secondary designers from Los Alamos and Lawrence Livermore and designers from Sandia participated. The workshop series, sponsored by ASCI and the National Science Foundation, brought together experts in high-performance computing, scientific visualization, emerging computer technologies, physics, chemistry, materials science, engineering, and nuclear weapons design. These experts worked to assess the needs of the scientific and engineering community, to identify current and projected capabilities, and to outline a federal research and development agenda in scientific visualization, human interfaces, and the manipulation of massive scientific data sets. This effort produced a report documenting the needs, most probable commercial technologies that would be available to address those needs, and areas of research to be undertaken. The report is titled “Data and Visualization Corridors, Report on the 1998 DVC Workshop Series” (CACR-164, September 1998).

Now on p. 41.
See comment 14.

The following are GAO's comments on the Department of Energy's letter, dated June 7, 1999.

GAO's Comments

1. DOE's comments relating to our recommendations and our responses are discussed in the executive summary, and chapters 2 and 3, where the recommendations appear.
2. The report text was revised to note that the Chiles Commission made this comment. The Commission was charged by the National Defense Authorization Acts of 1997 and 1998 to address the issue of maintaining nuclear weapons expertise.
3. We believe that the evidence presented in our report clearly demonstrates that long- and short-term planning efforts have been inconsistent and incomplete. For example, during the first 3 years of the program, annual implementation plans were prepared for some, but not all, components for only some of those years. In fact, in commenting on the draft of this report, DOE concurred with our recommendation that comprehensive planning and progress tracking systems were needed and cited improvements that they were making.
4. We agree that lower-level milestones can change for the reasons cited by DOE. However, the milestones discussed in this section of our report are the long-term, high-level program milestones. Furthermore, we believe that the examples cited in the report text and those shown in the chart in appendix I clearly demonstrate the many inconsistencies in DOE's strategic planning documents and the need to better document the many changes that have taken place.
5. The example referred to is just one of several cases in which DOE planning documents inconsistently reported the target completion date for a milestone. For example, the dates for the following milestones were reported differently in various documents: "Microaging," "3-D Casting Microstructure," "Macro/Micro Aging," "Prototype 3-D Physics," and "Component Deterioration Model". In addition, DOE provided us with many versions of DOE planning documents and briefings that contained similar inconsistencies. Furthermore, in some of these documents, the dates were reported in calendar years and in others in fiscal years. In still others, program and laboratory officials were uncertain of whether the dates were in calendar or fiscal years.

6. The DOE response to our referenced statement does not address the point of our statement. We state that laboratory milestones are reported inconsistently in the sources provided by DOE. The examples we use in our report illustrate this fact. DOE contends that transcription and printing errors caused labels and dates to change and that this problem “explains virtually all the contradictory information” about milestones detailed in the appendix to our report. We do not agree that these types of errors explain most of the problems we found in reviewing the documents provided to us. The point of our statement was that there is inconsistency among milestones found on the consolidated laboratory-level milestone chart and the two laboratory-level milestone charts used to support the consolidated chart. The examples we used show that milestones appear on the consolidated chart that are not found on the laboratory-level nuclear and nonnuclear milestone charts, which form the basis of the consolidated chart. Correspondingly, there are milestones on the laboratory-level nuclear and nonnuclear milestone charts that are not found on the consolidated chart.

DOE explains that these charts were not meant to provide details about milestones and that they are visual supplements to written documentation that provides detailed information about the milestones. This statement is surprising to us. One of the largest problems we had in evaluating this program stemmed from the fact that there was almost no detailed information about these milestones. As we state in our report, the only milestone for which we found a detailed description was for the primary burn code scheduled for completion in 1999. Our review of implementation plans, strategic plans, reports from principal investigators meetings, and other documents did not reveal any of the written documentation DOE claims to support the milestones found on these charts.

7. DOE and laboratory officials told us that the linkages referred to here are not obvious and that only those with technical expertise could understand them. At our request, program officials tried for over 4 months to prepare documents showing the linkages between projects and laboratory-level and higher-level milestones. Subsequently, program and laboratory officials told us that developing this information was a very useful exercise for them because the linkages were not always evident. DOE’s comments about the principal investigators’ meetings are not relevant to this point. However, as noted in another part of the report, there are few detailed reporting requirements for presentations at these meetings. In fact, we found that many of the accomplishments reported for individual projects

at the meetings were generally not tied back to the annual milestones established for those projects. In addition, there was no systematic assessment of how the results on individual projects are contributing to meeting either the laboratories' annual or program milestones.

8. We understand that project-level milestones and program-level milestones are not the same. However, we believe that DOE and the laboratories should be able to demonstrate how individual projects funded under the ASCI effort contribute to achieving annual laboratory-level and overall program-level milestones. This is essentially the "link" that we expected to see but did not find in program plans and related documents. In fact, in some cases program officials were unable to explain what that link was, although they stated that the projects were needed for the overall stockpile stewardship program.

9. DOE explained that they employ multiple research efforts and risk mitigation strategies to achieve program goals. We modified the report to clarify the use of multiple research approaches in the area of infrastructure development. We noted during our review, however, that these research efforts and risk mitigation strategies were poorly documented. DOE also explained the need to eventually select the most promising of these research efforts. However, DOE's claim that there are no set timetables for doing this seems inconsistent with the need to maintain the accelerated pace of the program. As stated in our report, program-planning documents should clearly identify research strategies, critical paths, and decision points.

10. We have added wording to the body of the report to recognize the completion of DOE's review of the nonnuclear mechanics codes. However, as stated in the report, at this juncture, most program-level milestones remain undefined.

11. As noted in our report, there are few detailed reporting requirements for presentations at these meetings. In fact, we found that the accomplishments reported for individual projects at the meetings were generally not tied to the annual milestones established for those projects. In addition, there is no systematic assessment of how the results on individual projects are contributing to meeting either the laboratories' annual or program milestones.

12. We do not believe that we misinterpreted the information provided to us. Documentation provided to us by Sandia officials and the statements

they made during our meeting with them on March 15, 1999, showed that the Sierra software development effort was 6 months or more behind schedule and that this slippage was affecting other software development efforts. Specifically, they cited the Fuego software development effort, which supports the “Abnormal Environment Thermal Assessment” milestone. They acknowledged that if any milestone in fiscal year 1999 is missed, it would most likely be this one because of the delays in Sierra. Sandia officials also stated that they had to redirect the Fuego, Coyote, Jas, and Pronto software development efforts until Sierra is ready. Whether the Sierra effort is 6 months behind schedule or whether the schedule itself has been slipped or, in DOE parlance, “rebaselined” by 6 months to recognize this slippage is irrelevant to the fact that the program schedule has been altered by delays in the Sierra project and that those delays are affecting other software efforts.

13. We believe that the statement in the report is correct as stands. We agree that implementation plans have improved over time. With respect to the quarterly reporting requirement for fiscal year 1999, it was added halfway through the fiscal year following our requests for information about progress towards meeting program milestones. We believe that this new quarterly reporting requirement is a positive improvement in the program’s progress tracking system.

14. DOE provided additional information that did not require any response or changes to the report.

15. The actual sustained speed achieved on the sPPM code was 800 billion operations per second or one-half that reported on the Linpack code. While we recognize that performance on the sPPM was a “best effort” in the contract, we believe that it is misleading to repeatedly cite the higher performance on the Linpack code without recognizing the fact that the computer has not yet passed its contractually required acceptance test. With respect to our statement that the Linpack code was substantially easier than the sPPM code, our basis was comments made to us by program officials at a meeting on September 3, 1998. At that meeting, program officials stated that they were using the sPPM code instead of the Linpack scale because the Linpack is a “toy program” not useful at all for measuring the capabilities needed by the weapons program. For example, Linpack is not an industry standard, contains only a few lines of code, and does not measure important capabilities such as the use of input/output devices. On the other hand, sPPM, using hydrodynamic calculations, fully tests the platform regarding communications devices, as well as how

efficiently the software scales to 6144 processors, and the ability to incorporate I/O devices.

16. Wording was added to the report to clarify that, according to the ASCI Program Plan, a 10-fold increase was expected from improving software to take advantage of parallel processing techniques.

17. We revised the report text to clarify that we found little documentation of software development progress as measured against established milestones.

18. As pointed out in comment 12, the Sierra project is 6 months or more behind schedule, is affecting other software development efforts, and could potentially result in a milestone being missed during fiscal year 1999. Simply changing the baseline used to measure progress does not change these facts.

19. We were provided with this information at a meeting with senior program officials at Sandia on March 15, 1999. Also, see comment 12.

20. DOE's comment is discussed in the agency comments section of chapter

21. The National Academy of Sciences information, as referenced by DOE, does not address the tracking of costs on research programs and thus does not support DOE's management approach as DOE contends.

22. The report does recognize that the ASCI Red computer at Sandia is capable of operating at a theoretical peak speed of 1.8 trillion operations per second. While this is a significant accomplishment, DOE fails to mention that the ASCI Red computer was developed by a different vendor and uses an architecture that is different from the architecture in the computer being developed at Los Alamos.

23. Although things may have changed since 1998 as DOE states, our point is still valid. Operating each new system constructed (including the new 3 trillion operations per second system at Los Alamos) as a fully integrated system poses challenges for the ASCI program. For example, ASCI's fiscal year 1999 implementation plan states, "The most critical issues affecting the successful implementation of the [Blue Mountain] system include hardware reliability and stability." Recognizing this fact, the contract statement of work for the Los Alamos computer included requirements to measure the stability of the system. However, as of March 1999, the

computer had not met these requirements, which clearly demonstrates this point.

24. We have revised the text to more accurately cite the Blue Ribbon Panel report.

25. We have revised the report text to more fully describe this survey effort.

Ordering Information

The first copy of each GAO report and testimony is free. Additional copies are \$2 each. Orders should be sent to the following address, accompanied by a check or money order made out to the Superintendent of Documents, when necessary. VISA and MasterCard credit cards are accepted, also. Orders for 100 or more copies to be mailed to a single address are discounted 25 percent.

Orders by mail:

U.S. General Accounting Office
P.O. Box 37050
Washington, DC 20013

or visit:

Room 1100
700 4th St. NW (corner of 4th and G Sts. NW)
U.S. General Accounting Office
Washington, DC

Orders may also be placed by calling (202) 512-6000 or by using fax number (202) 512-6061, or TDD (202) 512-2537.

Each day, GAO issues a list of newly available reports and testimony. To receive facsimile copies of the daily list or any list from the past 30 days, please call (202) 512-6000 using a touchtone phone. A recorded menu will provide information on how to obtain these lists.

For information on how to access GAO reports on the INTERNET, send an e-mail message with "info" in the body to:

info@www.gao.gov

or visit GAO's World Wide Web Home Page at:

<http://www.gao.gov>

**United States
General Accounting Office
Washington, D.C. 20548-0001**

**Bulk Rate
Postage & Fees Paid
GAO
Permit No. G100**

**Official Business
Penalty for Private Use \$300**

Address Correction Requested

