United States General Accounting Office

GAO

Report to the Chairman, Committee on the Budget, House of Representatives

July 1998

INFORMATION TECHNOLOGY

Department of Energy Does Not Effectively Manage Its Supercomputers

GAO/RCED-98-208

GAO

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

Resources, Community, and
Economic Development Division

B-280167

July 17, 1998

The Honorable John R. Kasich
Chairman, Committee on the Budget
House of Representatives

Dear Mr. Chairman:

As requested, this report examines the Department of Energy's (DOE) acquisition and use of supercomputers. Specifically, it (1) identifies the number and cost of the supercomputers DOE acquired in fiscal years 1994 through 1997 and the number and proposed funding for planned major supercomputer acquisitions in fiscal years 1998 through 2000; (2) determines the stated need for DOE's supercomputers, the utilization rates for them, and the potential for facilities to share these resources; and (3) identifies and describes the process DOE and its contractors employ to validate the need for additional supercomputers and contrasts that process with the technology investment process set forth in the Clinger-Cohen Act.

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

Please call me on (202) 512-7106 if you or your staff have any questions. Major contributors to this report are listed in appendix II.

Sincerely yours,

Susan D. Kladiva

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

Executive Summary

Purpose

Seven Department of Energy (DOE) national laboratories spent over $800 million acquiring and operating supercomputers—the largest and fastest computers currently being built—from fiscal years 1994 through 1997.[1] Over the next 3 years, DOE plans to acquire still more powerful and expensive systems to use for defense and energy research projects that require complex modeling, simulation, and computation.

Concerned about the status of DOE's efforts to acquire supercomputers, the Chairman, House Committee on the Budget, asked GAO to (1) identify the number and cost of the supercomputers DOE acquired in fiscal years 1994 through 1997 and the number and proposed funding for planned major supercomputer acquisitions in fiscal years 1998 through 2000; (2) determine the stated need for DOE's supercomputers, the utilization rates for them, and the potential for facilities to share these resources; and (3) identify and describe the process DOE and its contractors employ to validate the need for additional supercomputers and contrast that process with the technology investment process set forth in the Clinger-Cohen Act of 1996.

Background

DOE has been a world leader in computing for decades. Currently, a DOE-funded computer called "ASCI Red" at the Sandia National Laboratory holds the world's record for processing speed—1.8 trillion operations per second.[2] DOE plans to acquire much more powerful supercomputers with processing speeds of up to 100 trillion operations per second in the next several years. Such speeds are being achieved as a result of a new system configuration—called scalable massively parallel processing—that permits hundreds to thousands of processors to be linked together to achieve ever more powerful performance.[3]

In 1996, the Congress passed the Clinger-Cohen Act, which requires federal agencies to adopt a comprehensive approach to acquiring and managing information technology (including supercomputers), and

[1]There is no standard definition of a supercomputer. Furthermore, because of changing technology, a machine considered to be a supercomputer in 1993, for example, might not be considered one in 1998. For the purpose of this report, GAO used the TOP 500 Supercomputers List prepared by the University of Mannheim (Germany) and by the University of Tennessee as the basis for determining if a computer could be considered a supercomputer during the year it was acquired. In addition, GAO used this list for ranking DOE's supercomputers and determining DOE's share of total world capacity.

[2]Operations per second is a measure of how fast a computer retrieves and executes instructions (or operations) from memory. "Mega-," "giga-," and "tera-flops" are commonly used as shorthand terms for million, billion, and trillion operations per second, respectively.

[3]Processors are computer chips that essentially represent the "brains" of the computer.

charged the Office of Management and Budget with oversight responsibility.

Results in Brief

The Department of Energy has about 17 percent of the world's supercomputing capacity and is planning to almost triple its capacity over the next 3 years. Seven Department of Energy national laboratories have 42 supercomputers, 35 of which were acquired during fiscal years 1994 through 1997, at a cost of about $300 million, by the Offices of Defense Programs and of Energy Research. The cost of operating supercomputers is also substantial, totaling $526 million for the seven national laboratories for fiscal years 1994 through 1997. During fiscal years 1998 through 2000, the Department plans to acquire five major supercomputers at an estimated cost of $257 million.

Overall, the Department of Energy's national laboratories used only about 59 percent of their available supercomputing capacity in 1997[4] and are missing opportunities to share these resources.[5] Utilization rates varied among the laboratories from about 31 percent to about 75 percent. The sharing of supercomputers among Department of Energy laboratories and with DOE-funded off-site users is not generally considered as a way to better use existing resources and/or to forgo the need to acquire more supercomputers. In addition, the largest supercomputers are not being used to run the very large-scale programs that were used to justify their acquisition. In 1997, for example, less than 5 percent of the jobs run on the largest supercomputers used more than one-half of the machines' capabilities.

The Department of Energy has not effectively overseen the acquisition and use of supercomputers. Because the Department does not manage supercomputers as an agency resource, no person or office within the Department of Energy knows at a given time how many supercomputers there are, what they cost, or how they are being utilized. The Department lacks an investment strategy and a defined process to ensure that supercomputer acquisitions are fully justified and represent the best use of funds among competing priorities. Instead, the Department's existing management processes separate supercomputer acquisitions from the

[4]Although we cite 1997 utilization data throughout this report because they are the most current, we also obtained utilization data for fiscal years 1994 through 1997 that showed similar results.

[5]Utilization data—which measure the percent of available time that a system's processors are in use—can be automatically generated by most supercomputers but are not systematically kept by the laboratories or the Department of Energy.

projects they support, and the Department's chief information officer does not oversee the acquisition or use of supercomputers. As a result, new supercomputers are planned and acquired with little departmental oversight, while underutilized capacity already exists within the Department. The Department of Energy's proposed implementation of the Clinger-Cohen Act will not improve departmental oversight. The Clinger-Cohen Act requires that federal agencies implement a comprehensive, efficient approach to acquiring and managing information technology. In April 1998, the Department outlined its plan to implement a new investment planning and oversight process for information technology in response to the Clinger-Cohen Act. DOE's new process separately manages administrative and scientific computers, leaving the responsibility for scientific computers—including supercomputers—to individual program offices. This proposed approach reflects the view of the Department's program offices that supercomputers are research "tools" rather than information technology investments. This approach may also allow DOE's program offices to continue acquiring supercomputers outside the Department's normal process for complying with the Clinger-Cohen Act. Contrary to what is envisioned in the Clinger-Cohen Act, this approach effectively places the vast majority of DOE's information technology resources outside the purview of the Department's chief information officer.

In addition, the cost and significance of the supercomputers being developed under the Department's Accelerated Strategic Computing Initiative (ASCI) warrants that program's being designated as a "strategic system" subject to the highest level of departmental oversight. The ASCI program was created in response to a presidential decision in 1995 to sign the Comprehensive Test Ban Treaty. The program is designed to provide the unprecedented simulation capabilities needed to help verify the safety and reliability of U.S. nuclear weapons without nuclear testing. It also has the ambitious goal of increasing computer performance at a rate of development exceeding the current and projected state of the art, with the ultimate goal of developing a supercomputer capable of 100 trillion operations per second by 2004. The ASCI program is estimated to cost about $4 billion from fiscal years 1996 through 2010. Funding for the ASCI program is $223.5 million in fiscal year 1998.

Principal Findings

DOE's Supercomputer Capabilities Have Grown Significantly

DOE's supercomputing capabilities have grown over 10-fold from fiscal years 1994 through 1997. Seven DOE national laboratories have 42 supercomputers, 35 of which were acquired over this period at a cost of about $300 million. The average cost of the supercomputers acquired since fiscal year 1994 is about $8.5 million, but some large-scale supercomputers have cost several times as much. For example, in 1997, a supercomputer at Lawrence Berkeley cost over $25 million, and the record-holding ASCI Red system cost about $55 million. The Department's stated need for supercomputers is to perform increasingly complex computations and modeling to support two main activities: (1) maintaining the safety and reliability of U.S. nuclear weapons under the Department's stockpile stewardship program, which is administered by the Office of Defense Programs, and (2) conducting civilian research into complex scientific problems—such as global climate change, human gene structure, and environmental contamination—in projects administered by the Office of Energy Research.

The costs of operating a supercomputing facility are substantial and can equal or exceed the acquisition costs: For fiscal years 1994 through 1997, the seven national laboratories spent about $526 million to operate their supercomputers. At a large laboratory like Los Alamos, operating costs were over $56 million in fiscal year 1997. Costs vary depending on the size of the facility, the type of research being conducted, and the cost of associated equipment.

DOE is planning to acquire more expensive systems in the future. GAO identified five major supercomputer acquisitions planned for fiscal years 1998 through 2000, including planned increments to existing large systems at four laboratories, that are budgeted to cost $257 million. Two of these systems—ASCI Blue Mountain at Los Alamos and ASCI Blue Pacific at Livermore—will cost $77.8 million and $54.4 million, respectively, during these 3 years.

DOE's Supercomputers Are Underutilized, and Opportunities to Share Them Are Missed

Overall, DOE's national laboratories used about 59 percent of their available supercomputing capacity in 1997. Nevertheless, DOE, which has about 17 percent of the world's total available supercomputing capacity, is planning to almost triple its capacity over the next 3 years. Utilization rates among individual laboratories varied from about 31 percent to about 75 percent. Utilization rates for individual systems and laboratories may vary because of a number of factors, including the age and type of machine (whether it is an experimental machine or a more stable production model) and the use of the machine. For example, Lawrence Berkeley which had the highest utilization rate, 75 percent, is a designated user facility, available to DOE-funded researchers across the country. However, the largest supercomputers—those being justified as needed to run very large computer programs simultaneously across hundreds or thousands of processors to solve the largest problems in a reasonable time—are seriously underutilized. GAO found that during fiscal year 1997, less than 5 percent of the jobs run on the largest and fastest supercomputers at DOE's laboratories used more than one-half of the supercomputer's available processors.

With the exception of Lawrence Berkeley, there is only limited sharing of supercomputers among DOE laboratories and with DOE-funded off-site users. At the end of 1997, about 41 percent of DOE's total supercomputer capacity (equal to about 1.7 trillion operations per second) was not being used. Nevertheless, DOE plans to increase its capacity by another 1 trillion operations per second during 1998. At the same time, DOE is missing opportunities to share its unused capacity among laboratories because it does not emphasize looking for such opportunities. In previous years, DOE has collected utilization data on laboratories' computers and required that computer managers analyze workload data and consider sharing opportunities prior to acquiring new computers. However, those requirements were canceled in September 1995 as part of an effort to reduce burdens on contractors and eliminate paperwork. DOE no longer emphasizes the consideration of utilization rates and sharing opportunities when laboratories are seeking funding for a new supercomputer.

DOE Has Not Been Effectively Overseeing Supercomputer Acquisitions

DOE does not have a process to ensure that supercomputer acquisitions are fully justified and represent the best use of funds among competing priorities. Overall, DOE lacks an investment strategy for acquiring supercomputers and does not follow the Clinger-Cohen Act's criteria, which require that DOE and other federal agencies implement a comprehensive, efficient approach to acquiring and managing information technology.

In April 1998, DOE outlined its new planning and oversight process, under the Clinger-Cohen Act, for investing in information technology. The new process separately manages administrative and scientific computers—including supercomputers—that DOE's program offices view not as information technology acquisitions but as "tools" supporting research projects. As envisioned, the "dual track" approach leaves decisions on supercomputers to individual program offices and may allow program offices to continue acquiring supercomputers outside the Department's normal Clinger-Cohen Act process. This approach, contrary to what is envisioned in the Clinger-Cohen Act, also places the vast majority of DOE's information technology resources outside the purview of the Department's chief information officer.

DOE has established criteria for designating projects costing over $400 million that are an urgent national priority, are high risk, have international implications, or are vital to national security as "strategic systems." The purpose of designating strategic systems is to ensure informed, objective, and well-documented decisions about key events for such systems, such as changes to baseline costs and schedules. Among DOE's 11 current strategic systems are the National Ignition Facility and the Tritium Supply Facility, both of which, like the ASCI program, are related to ensuring the reliability and safety of the nuclear weapons stockpile. In prior years, the Department has not effectively managed such acquisitions, and they have often been late and over budget. The ASCI program will likely cost about $4 billion for fiscal years 1996 through 2010, is an urgent national priority because of national security concerns, and has international implications because it is a major factor in the United States' support of a Comprehensive Test Ban Treaty. In addition, the program is high risk because it seeks to advance the state of the art in supercomputing and simulation well beyond current capabilities, has already experienced delays and seen its projected cost increase, and depends on as yet unknown technologies for success. These factors warrant designating the ASCI program as a "strategic system" to avoid the types of problems historically encountered by DOE with such projects. However, the program has not been so designated.

Recommendations

GAO is making several recommendations to improve the utilization and sharing of supercomputers at DOE's laboratories. Specifically, GAO recommends that the Secretary of Energy adopt a Clinger-Cohen approach for acquiring supercomputers that (1) pertains to all Department-funded supercomputers; (2) requires the consideration of utilization data and sharing opportunities prior to funding new supercomputer acquisitions; and (3) includes sufficient justification to allow for the meaningful consideration of alternatives. Finally, GAO is recommending that the Secretary designate the ASCI program as a “strategic system” warranting departmental oversight at the highest level.

Agency Comments

GAO provided a copy of a draft of this report to the Department of Energy for review and comment. DOE generally disagreed with the findings and recommendations in the report.

DOE stated that measuring processor utilization takes in only one dimension of massively parallel computers and does not account for other factors, such as memory size, that can affect utilization. As a result, according to DOE, in the past supercomputers became saturated at about the 70-percent utilization level, and today’s generation of massively parallel computers becomes saturated at a substantially smaller percentage. DOE stated that it would be impossible as a result to achieve 100-percent utilization using available processor time as the measure. Furthermore, DOE stated that its utilization data and experience from the past 30 years indicate that the highest possible level of utilization would be considerably less than 100 percent.

GAO agrees that available processor time is not a perfect measure of supercomputer use; however, it is the only one currently in widespread use that can be applied to a number of different computer architectures and models. Furthermore, neither DOE nor anyone else GAO spoke with was able to provide a better measure than processor utilization. With respect to what constitutes an acceptable utilization rate, GAO did not suggest that DOE should, or even could, achieve 100-percent use. Rather, the draft report stated that DOE’s overall utilization rate of 59 percent was low compared with the higher rates of 70 to 75 percent achieved at some of the Department’s laboratories and on some individual computers. The draft report also stated that DOE is missing opportunities to improve utilization because it does not monitor utilization or seek out sharing opportunities. Given that DOE is able to obtain these higher utilization rates on some computers and at some sites, GAO continues to believe that DOE could achieve better levels of utilization on many of its supercomputers.

DOE also commented on an example used in the report of the types of programs that were being run on the largest supercomputers, whose acquisition DOE had justified as needed because of their capability for running very large-scale programs on hundreds or thousands of processors. The report states that less than 5 percent of the programs being run used one-half or more of the supercomputer's available processors. DOE stated that the 5 percent of the programs run on the ASCI Red supercomputer at Sandia accounted for over 80 percent of the utilization of that supercomputer. GAO believes that while DOE's statement is true, it is misleading. The utilization rate on the ASCI Red supercomputer at Sandia is about 43 percent in total. And about 80 percent of that rate (or about 34 percent of the total available time) is accounted for by running a few very large-scale jobs that take a considerable amount of time. Consequently, up to 57 percent of the total processor time available on this supercomputer still could be used for other jobs.

DOE also stated that GAO's conclusion that 41 percent of its overall supercomputer capacity is available for sharing was erroneous. DOE based its statement on the 80-percent figure DOE cited for the Sandia ASCI supercomputer and concluded that the true percentage of its unused capacity is close to zero. DOE also stated that the sharing of the ASCI program supercomputers is very difficult because of national security concerns. GAO disagrees. In fact, as GAO pointed out above, the total use on the ASCI Red supercomputer, including the very large programs, is only 43 percent, and up to 57 percent is still available for other use. In connection with the sharing of the ASCI supercomputers, they were originally planned and are being installed to allow just this type of sharing. To date, the three ASCI supercomputers have been designed to have both classified and unclassified parts that can also, after following proper procedures, be linked together to run the largest programs. In fact, one of the requirements of the Sandia ASCI supercomputer was that it could be switched between classified and unclassified uses in less than 30 minutes. In addition, ASCI program documents state that 10 percent of the capacity of these supercomputers will be available to users from outside DOE's laboratories, such as the universities participating in research for the ASCI program.

DOE disagreed with GAO's finding that it lacks an investment strategy for acquiring supercomputers and that its proposed implementation of the Clinger-Cohen Act does not meet the act's criteria and GAO's recommendation that DOE should adopt a Clinger-Cohen process that pertains to all DOE-funded computers. As stated in the report, DOE believes that scientific computing should be treated differently from other computer resources within the Department and should be subject to review by either the Offices of Energy Research or Defense Programs rather than by the departmentwide Executive Committee on Information Management. GAO believes these offices' acquisitions of scientific information technology should not be excluded from departmentwide oversight. In fact, since DOE contractors, including the laboratories that acquire supercomputers funded by program offices, account for over 80 percent of the Department's spending on information technology, GAO believes it is crucial that these program offices' supercomputer acquisitions be included in a unified DOE-wide approach to implementing the Clinger-Cohen Act. To do otherwise effectively puts the Department's most valuable information technology outside the purview of the agency's chief information officer and the Executive Committee on Information Technology. Accordingly, GAO stands behind its recommendation that DOE should adopt a departmentwide process that meets the Clinger-Cohen Act's criteria and includes supercomputers and other scientific computing resources.

Finally, DOE disagrees with GAO's recommendation that the ASCI program should be designated as a strategic system, saying that the program already gets high-level oversight as part of the Clinger-Cohen Act process and through normal departmental channels. GAO disagrees. First, as cited in the report, the ASCI program is critical to the efforts to ensure the safety and reliability of the nation's stockpile of nuclear weapons and meets all the other criteria for being designated as a strategic system. Second, the process that DOE is implementing in response to the Clinger-Cohen Act would allow the same program office that has a vested interest in the program as its manager to serve as the oversight body for the Department's supercomputers. GAO continues to believe that this approach does not follow the Clinger-Cohen Act nor achieve the degree of high-level oversight that designation as a strategic system would provide. Given that DOE's process does not follow the Clinger-Cohen Act criteria and the importance of the ASCI program to ensuring the safety and reliability of the nation's nuclear weapons stockpile, GAO continues to recommend that the Secretary should designate the ASCI program as a strategic system. DOE's comments and GAO's responses are in chapters 3 and 4 and appendix I.

Contents

Figures

Abbreviations

ASCI	Accelerated Strategic Computing Initiative
DOE	Department of Energy
GAO	General Accounting Office
OMB	Office of Management and Budget
R&D	research and development

Introduction

A supercomputing revolution is under way in the United States as dramatic advances in supercomputers continue—doubling in power approximately every 18 months. Supercomputers are defined as the largest and fastest computers currently being built. Today's supercomputers are capable of processing speeds of up to 1,000 times greater than they were capable of 5 years ago. Federal agencies, including the Department of Energy (DOE), have been at the forefront of this revolution. A DOE-funded supercomputer at the Sandia National Laboratory that is capable of about 1.8 trillion operations per second holds the world record for processing power. DOE, working with industry, has plans to build even faster computers, with a goal of 3 trillion operations per second by mid-1998, and 100 trillion by 2004.

DOE's Use of Supercomputers

Currently, seven DOE national laboratories and two DOE atomic power laboratories have supercomputers. DOE's program offices fund supercomputer purchases (or leases), and the laboratories' management and operating contractors acquire and operate the systems. According to DOE, new supercomputer acquisitions are relatively frequent because of rapidly changing technology. Since 1993, when statistics were first systematically collected, DOE has consistently had several supercomputers that have ranked among the world's most powerful as measured by a list of the top 500 supercomputers in the world. As of November 1997, for example, DOE had the first and fifth most powerful computers in the world and 10 of the top 100 supercomputers. Overall, DOE's laboratory supercomputers accounted for about 17 percent of the total supercomputer capacity in the world. DOE's supercomputing capability has grown over 10-fold from fiscal year 1994 through fiscal year 1997, from a total for DOE laboratories of about 360 billion operations per second to about 3.8 trillion operations per second. Under DOE's plans, this total capability will increase another 280 percent by fiscal year 2000, as shown in figure 1.1.

Figure 1.1: Trends in DOE's Supercomputing Capability, Fiscal Years 1994-2000

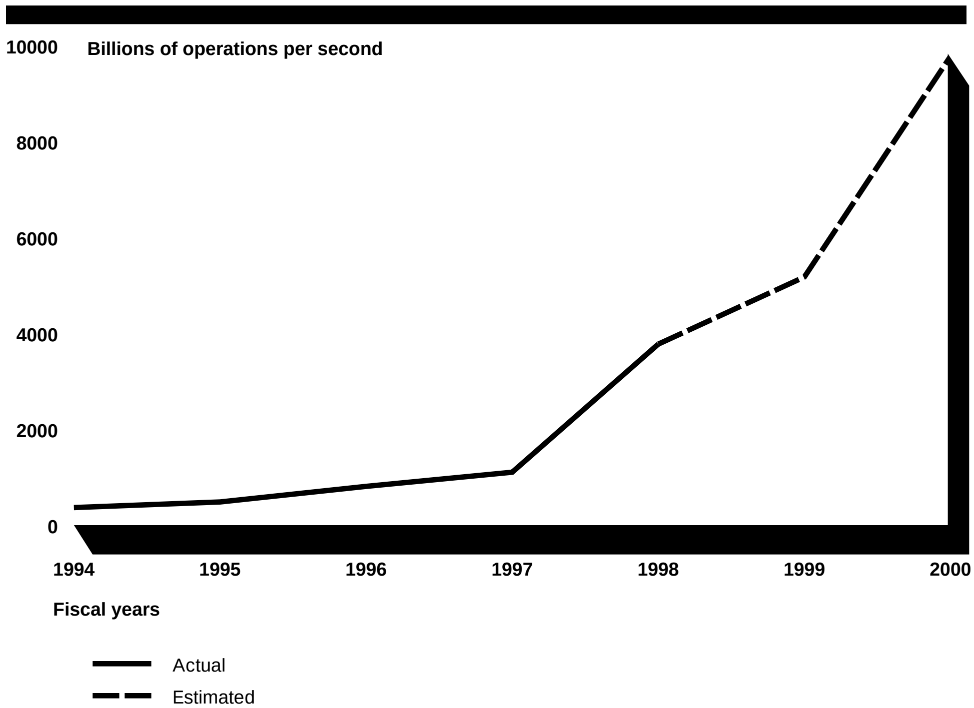

Source: TOP 500 Supercomputers List and DOE.

Supercomputers enable DOE to do leading-edge scientific research. Advances in computing power have been made possible by a new supercomputer configuration (or "architecture") called massively parallel processing. Past supercomputers (with "sequential" or "vector" architectures) executed operations more or less sequentially, using at most a few processors working concurrently. The latest supercomputers, however, execute many operations in parallel, simultaneously using hundreds, and even thousands, of processors. Such systems are referred to as scalable, massively parallel systems. In experimenting with such systems, researchers have linked more and more processors together to achieve ever greater processing speeds. To accommodate these newer, larger systems, concurrent research in software programs, interconnections, and graphics capabilities have been necessary.

As the year 2000 approaches and computing operations on the scale of a trillion operations per second become more routine, DOE is funding a variety of research initiatives to take advantage of the new simulation and computational capabilities of supercomputers. These initiatives include increasingly complex computations and the simulation of nuclear weapons, the global climate, the environment, pollution, and human gene structure. DOE uses supercomputers to support two major research missions: (1) ensuring the safety and reliability of nuclear weapons, under its Office of Defense Programs' nuclear stockpile stewardship program and (2) attempting to solve nondefense science and engineering problems, called "Grand Challenges," under its Office of Energy Research.

The Office of Defense Programs funds nuclear stockpile computing efforts to simulate the behavior of nuclear weapons. The Accelerated Strategic Computing Initiative (ASCI) was created by DOE in response to a presidential decision in 1995 to sign the Comprehensive Test Ban Treaty. ASCI is designed to provide the unprecedented simulation capabilities needed to help verify the safety and reliability of U.S. nuclear weapons without nuclear testing. ASCI has the ambitious goal of achieving increasing computer speeds at a rate of development exceeding the current and projected state of the art in coming years. Speeds of 3, 10, 30, and ultimately 100 trillion operations per second by 2004 are envisioned. ASCI funding for fiscal year 1998 was $223.5 million. ASCI-related funds are also contained in DOE's stockpile computing budget. On the other hand, the Office of Energy Research funds nondefense computational research projects, including specific grand challenges that require large-scale supercomputing capability. Grand challenge research is coordinated on an interagency basis through the High Performance Computing and Communications (HPCC) program, a $1-billion-a-year supercomputing initiative stemming from the High Performance Computing Act of 1991. This act was intended to accelerate the development of advanced technologies for the information age. DOE's fiscal year 1997 budget request for the HPCC program was $124.6 million, and its request for grand challenges research funding was about $10 million. According to DOE, about 80 percent of the $1.4 billion it spends each year for information management is spent by the management and operating contractors that run its major facilities, including the national laboratories.

Figures 1.2 through 1.4 show the types of supercomputers in DOE's national laboratories.

Figure 1.2: A 512-Processor Cray T3E Supercomputer at Lawrence Berkeley, With a Capacity of 460 Billion Operations Per Second, Acquired in Fiscal Year 1997

Source: DOE.

Figure 1.3: An 8-Processor Cray YMP Supercomputer at Los Alamos, With a Capacity of 2.7 Billion Operations Per Second, Acquired in Fiscal Year 1990

Figure 1.4: The 9,168-Processor "ASCI Red" Supercomputer at Sandia, With a Capacity of 1.8 Trillion Operations Per Second, Acquired in Fiscal Year 1997

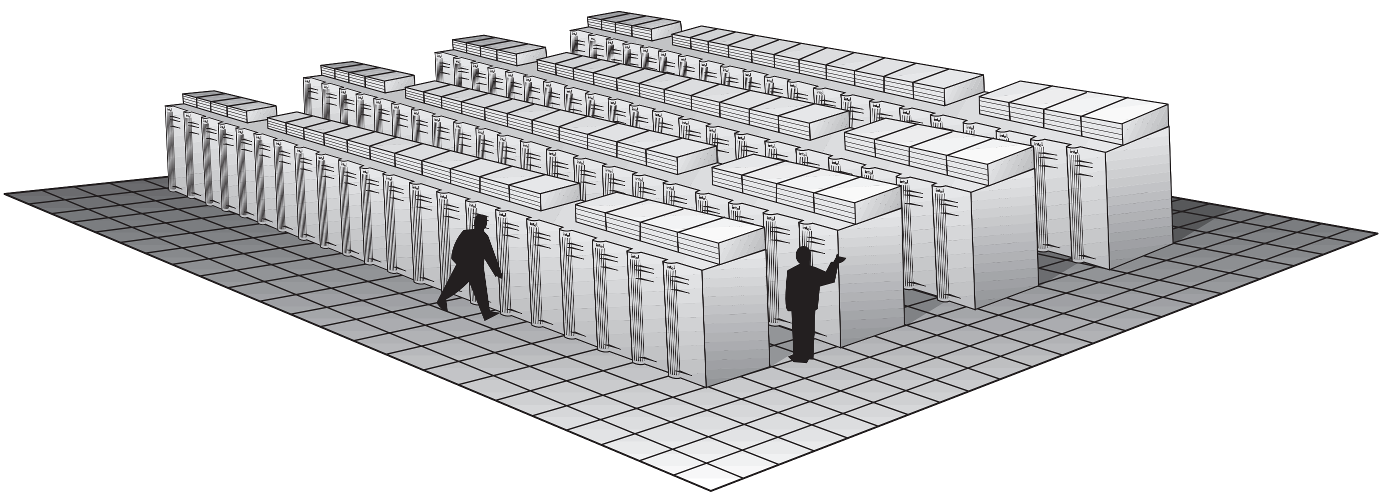

Source: Both figures provided by DOE.

Legislation Has Addressed the Acquisition of Computers

The Congress has shown an ongoing interest in supercomputers and information technology, both from the perspective of helping to ensure U.S. leadership in the field—as with the High Performance Computing and Communications Act of 1991—by promoting the efficient acquisition and management of computers. Under the Brooks Act of 1965, the Congress gave the General Services Administration the central authority within the federal government for acquiring information technology. In 1996, the Congress passed the Clinger-Cohen Act, which repealed the Brooks Act and gave the heads of agencies the authority to procure information technology directly. The act requires agencies to adopt a comprehensive approach to acquiring and managing information technology (including supercomputers) and charges the Office of Management and Budget with oversight responsibility.

Objectives, Scope, and Methodology

As requested by the Chairman, House Committee on the Budget, we (1) identified the number and cost of the supercomputers DOE acquired in fiscal years 1994 through 1997 and the number and proposed funding for planned major supercomputer acquisitions in fiscal years 1998 through 2000; (2) determined the stated need for DOE's supercomputers, the utilization rates for them, and the potential for facilities to share these resources; and (3) identified and described the process DOE and its contractors employ to validate the need for additional supercomputers and compared that process with the technology investment process set forth in the Clinger-Cohen Act. The scope of this review was departmentwide, including all facilities having supercomputers, although our principal focus was on DOE's multiprogram national laboratories. We also gathered information on the number and cost of supercomputers at two single-program laboratories—the Bettis and Knolls Atomic Power Laboratories. However, we did not perform detailed work at these two laboratories because of time constraints, the narrow focus and unique mission of their research programs, and their distinctive program management under the Office of Naval Reactors within DOE.

To identify the number and cost of DOE's existing supercomputers, we requested that DOE's Offices of Defense Programs and Energy Research provide us, for each supercomputer that they had funded, the year of purchase or lease and the cost (including the costs of any major upgrades after the initial purchase of the computer). The Office of Energy Research provided information on supercomputers at the Argonne, Lawrence Berkeley, Los Alamos, Oak Ridge, and Pacific Northwest National Laboratories. The Office of Defense Programs provided information on supercomputers at the Lawrence Livermore, Los Alamos, and Sandia National Laboratories. The Office of Naval Reactors provided information on supercomputers at the Bettis and Knolls Atomic Power Laboratories. We visited the seven national laboratories listed above and verified the data provided, although we did not independently verify the costs reported. In most cases, we relied on the cost data provided by the program offices, while in several other cases, we obtained clarifying cost information during our laboratory visits. To identify major planned supercomputer acquisitions, we relied primarily on information provided by the Offices of Energy Research and Defense Programs. Where possible, we sought to verify this information either during visits to the relevant laboratories or through documents used in formulating the agency's budget requests.

Recognizing that the initial purchase or lease cost of a supercomputer may represent only a portion of the cost of operating a supercomputer, we also developed information on the total costs of operating supercomputer centers at the seven national laboratories. To do this, we requested that the laboratories provide cost data for all of the cost accounts that constitute the supercomputing effort at the laboratory. We asked that these costs be divided into direct labor and "other" operating costs, and that they include any relevant program and laboratory overhead expenses. We requested this information for fiscal years 1994 through 1997 and asked that the laboratories project these costs for fiscal years 1998 and 1999. Where a cost center included activities unrelated to supercomputing, we asked the laboratory staff to estimate the proportion of costs attributable to the supercomputing effort. Finally, we asked that the laboratories reconcile the cost information provided with the information they supplied to DOE's budget and reporting system. We then provided draft summaries of this information to each of the laboratories and DOE's Office of the Chief Financial Officer for their review and comment or concurrence.

To determine the utilization rates for DOE's supercomputers, we held discussions with knowledgeable staff at Argonne and Sandia National Laboratories to learn how they collected and analyzed utilization data. We then developed a standardized data request for each supercomputer to determine its utilization rate on the basis of the number of processor hours available for running computer applications.[6] In addition to the processor

[6]The number of processor hours available for running computer applications is a function of the number of processors in the computer and the number of hours during the time period being considered. For example, a 256-processor supercomputer during a 30-day month has 184,320 processor hours available (256 processors x 30 days x 24 hours = 184,320).

hours available, we asked for the number of processor hours actually used for applications and the number of processor hours the computer was down for maintenance or repairs. We subtracted the number of hours that machines were down for maintenance or other reasons from the total number of hours available, to arrive at the number of hours the machines were actually available to users. We also wanted to determine the types of tasks the computers were used for because the largest, fastest, and most expensive supercomputers are being justified on the need for the capability to run the largest, most complex applications in a reasonable amount of time. Such jobs would be expected to use all or a large portion of the supercomputer's processors. Thus, for the newer supercomputers—those with over 128 processors—we asked the laboratories to provide information on the number of jobs using various ranges of available processors, that is 0 to 25 percent, 26 to 50 percent, 51 to 75 percent, and over 75 percent. We then analyzed the responses to determine whether the bigger supercomputers were being used to run the big computing jobs for which they were purchased (i.e., the percent of computing jobs that used over one-half of the available processors).

We took several steps to ensure that the data provided to us on the utilization of supercomputers by the laboratories were reliable. First, we visited the Lawrence Livermore, Los Alamos, and Sandia laboratories, which, taken together, have at least one supercomputer from all of the major manufacturers. We noted that collecting standard utilization data for each manufacturer and computer was possible because the utilization data are produced within the computers' operating systems using the same algorithms. During visits to these laboratories, we discussed controls over the systems, including physical access controls, which we found to be stringent because of the nature of the work performed by the laboratories. Each laboratory reviewed the data for reasonableness either periodically and/or before it sent us the data (e.g., checking to ensure that the utilization reported did not exceed total time available). We also sent a brief questionnaire to four other laboratories that provided utilization data and found similar processes and controls. On the basis of this review of the process, controls, and data reviews related to the creation of the utilization data provided us, we concluded that the utilization data are sufficiently reliable for use in this report.

To examine DOE's needs determination process for supercomputers, we interviewed responsible officials and obtained pertinent documentation at DOE's headquarters offices—Defense Programs, Energy Research, and the Office of the Chief Information Officer—as well as at selected DOE operations offices and at the various national laboratories where supercomputers are in use. From these sources, we obtained an overview of existing and past DOE procedures for justifying the acquisition of supercomputers, as well as a newly approved justification process intended to implement the Clinger-Cohen Act. Furthermore, we examined the details of DOE's informal and formal program and project planning and budgeting processes. DOE uses these processes to validate the need for projects that may also include supercomputers. We compared the information obtained to the requirements of the Clinger-Cohen Act. We also analyzed whether DOE's most expensive supercomputers, such as its ASCI computers, should be managed by the Department as strategic systems under its life-cycle asset management process.

We conducted our review from September 1997 through June 1998 in accordance with generally accepted government auditing standards.

Number, Capability, and Cost of DOE's Supercomputers Have Grown Rapidly

DOE has experienced rapid growth in the number, capability, and cost of its supercomputers since fiscal year 1994. As of December 31, 1997, DOE's laboratories had 42 supercomputers, several of which ranked among the most powerful in the world. From fiscal year 1994 through fiscal year 1997, DOE's supercomputer capability grew by more than 10-fold. Thirty-five of DOE's supercomputers were purchased in these years at a cost of about $300 million (in current year dollars). These supercomputers were funded either by the Office of Energy Research or the Office of Defense Programs.[1] Supercomputers are also expensive to maintain and operate. The overall cost of supercomputing, in addition to the costs of the supercomputers, for the seven DOE laboratories we examined was about $526 million for fiscal years 1994 through 1997.

DOE currently plans to spend an additional $257 million on major supercomputer acquisitions during fiscal year 1998 through fiscal year 2000. The newest, fastest supercomputers commercially available can cost $20 million or more to buy or lease. Some very large developmental models can cost more than $100 million. These acquisitions will increase DOE's supercomputing capability by almost 300 percent.

The Number, Capability, and Cost of DOE-funded Supercomputers Have Increased

DOE's seven national laboratories have significant supercomputer capacity, including several machines that are among the most powerful in the world. The laboratories had 42 supercomputers at the end of fiscal year 1997, up from 25 supercomputers at the start of fiscal year 1994. DOE's stated need for acquiring supercomputers is to perform increasingly complex computations and modeling to support two main program missions—(1) maintaining the safety and reliability of U.S. nuclear weapons under the Department's stockpile stewardship program, administered by the Office of Defense Programs, and (2) conducting civilian research into complex scientific problems, such as global climate change, human gene structure, and environmental contamination, in projects administered by the Office of Energy Research.

These acquisitions have significantly enhanced the capacity and capability of DOE's supercomputers. For example, at the start of fiscal year 1998,

[1]In addition to the 42 supercomputers located at the seven national laboratories, DOE's Office of Nuclear Energy reported having eight supercomputers, four of which were purchased in fiscal years 1994-97 at a cost of $22.4 million. These eight machines—obtained to support the Naval Reactors Program—are located at either Bettis or Knolls Atomic Power Laboratories. We are not including those machines among the 42 reported in this report because of their unique mission and because we did not visit these laboratories to verify the data provided by the Office of Nuclear Energy. These laboratories plan to replace some of their older supercomputers during fiscal years 1998 through 2000 at an estimated cost of $48 million.

DOE's supercomputers had the capacity to run about 3.8 trillion operations per second—more than 10 times the total capacity it had at the start of fiscal year 1994. In addition, individual machines, which can now have several dozen to several thousand processors linked together, are much more powerful. The newer supercomputers are capable of running software programs on several hundred to several thousand processors simultaneously to solve complex computer models. A DOE-funded supercomputer at Sandia which is capable of 1.8 trillion operations per second, is currently ranked as the fastest in the world.

DOE purchased 35 of the supercomputers in fiscal years 1994 through 1997 at a cost of about $300 million (in current year dollars). These supercomputers were funded by either the Office of Energy Research or the Office of Defense Programs. Table 2.1 shows the location and cost of the 35 supercomputers.

Table 2.1: Location and Cost of DOE-Funded Supercomputers

Current Year Dollars in Millions

National Laboratory	Number purchased, fiscal years 1994-97	Cost
Argonne	1	$16.7
Lawrence Berkeley	5	34.3
Los Alamos	5	75.1
Lawrence Livermore	10	79.1
Oak Ridge	3	13.4
Pacific Northwest	2	13.1
Sandia	9	69.8
Total	**35**	**$301.5**

Notes: Includes the cost of the initial purchase and subsequent upgrades. In some cases, the costs shown include items other than the supercomputer hardware. For example, some contracts contain provisions for maintaining the machine, writing software for the machine, or providing funds to assist in the development of the supercomputer or related critical technologies.

Source: Data provided by the Offices of Energy Research and Defense Programs and verified by GAO during site visits.

In addition to increasing in power, DOE's newest supercomputers are increasingly costly. While the average cost of a supercomputer acquired by DOE since fiscal year 1994 is about $8.5 million, the cost of the newer, more powerful supercomputers can be significantly more. For example, the commercially available Cray T3E-900 at Lawrence Berkeley, which was the fifth most powerful computer in the world in November 1997, cost over $25 million. In comparison, the largest machines being developed by

DOE, in conjunction with industry, can cost several times that amount. The ASCI Red supercomputer cost about $55 million in fiscal year 1997. The ASCI Blue Pacific supercomputer at Lawrence Livermore (3 trillion operations per second) will cost about $95 million, including development and some support costs. Los Alamos is acquiring a similarly sized machine, known as the ASCI Blue Mountain, at a cost of about $135 million. DOE's Office of Energy Research is funding a smaller version (1 trillion operations per second) of the ASCI Blue Mountain machine, also at Los Alamos, at a cost of about $40 million. In February 1998, the Office of Defense Programs funded another ASCI effort at Lawrence Livermore. This $85 million contract funds the next planned increment of the ASCI program for a computer capable of 10 trillion operations per second.

The costs of operating supercomputers are also substantial. These costs at the seven laboratories totaled about $526 million in fiscal years 1994 through 1997, excluding the costs of acquiring the supercomputers.[2] These costs varied depending on the size of the facility and the programs involved. At one of the larger facilities, Los Alamos, for example, the annual costs of supercomputing were over $56 million in fiscal year 1997. Supercomputers consume large amounts of electrical power and often require special, or additional, air conditioning equipment. For example, electricity costs for the computing center at Livermore are almost $1 million per year. In addition, software for supercomputers, especially those with massively parallel architectures, is generally not available commercially. As a result, additional resources must be available to help users develop, convert, or optimize their applications to run on these machines.

The ASCI program has contributed to significant increases in supercomputing costs at the three weapons laboratories since it started in fiscal year 1996. At Livermore, for example, the overall costs of supercomputing increased by about 30 percent from fiscal year 1995 to about $35.2 million in fiscal year 1997. Sandia is constructing a new building largely to support ASCI supercomputing, with a budgeted total cost of about $29 million between fiscal years 1999 and 2001. Similarly, Livermore is nearing completion of a $12 million renovation of an existing building to support its next generation of ASCI supercomputers. Most of the Livermore cost results from bringing in the electrical power to run and cool the supercomputers.

[2]The programs that house and operate these supercomputers impose various "burdens," or "overhead" charges, to the operating and equipment costs, as do the laboratories themselves. Those costs are also included.

Planned Major Supercomputer Acquisitions

DOE expects that its planned major supercomputer acquisitions for fiscal years 1998 through 2000 will cost about $257 million. These acquisitions represent the minimum that DOE plans to spend for supercomputers during those years because smaller acquisitions and upgrades are not included.[3] For several of these machines, including the ASCI-funded machines at Los Alamos and Lawrence Livermore and the Energy Research-funded machine at Los Alamos, the $257 million represents only the future funding increments for those machines that are being acquired and installed over more than one fiscal year. Table 2.2 shows DOE's planned major acquisitions for fiscal year 1998 through fiscal year 2000.

Table 2.2: Major Planned Supercomputer Acquisitions, Fiscal Years 1998-2000

Dollars in millions

Laboratory/ name of computer/status	Peak performance goal (trillions of operations / sec.)	Fiscal year 1998	Fiscal year 1999	Fiscal year 2000
Los Alamos/ASCI Blue Mountain/ongoing	3.1 (when completed)	$39.8	$38.0	$0
Livermore/ASCI Blue Pacific/ongoing	3.2 (when completed)	27.3	27.1	0
Los Alamos/Nirvana Blue/ongoing	1.0 (when completed)	8.0	8.0	8.0
Livermore/ASCI "Option White"/new start-Feb. 1998	10.0 (when completed)	21.3	28.3	35.5
Berkeley/new computer/first-year lease	[a]	0	6.7	9.0
Total	[a]	**$96.4**	**$108.1**	**$52.5**

[a]Lease details not final. The total cost is expected to be about $27 million.

Source: DOE.

These planned acquisitions will increase DOE's total supercomputing capability by another 280 percent by fiscal year 2000. Lawrence Livermore and Los Alamos are acquiring and installing supercomputers even larger

[3]In addition, as noted earlier, the Bettis and Knolls atomic power laboratories plan to replace some of their older supercomputers at a reported cost of about $48 million. This cost is not included in the totals reported in the body of this report.

than the ASCI Red machine—each capable of over 3 trillion operations per second. The "Option White" supercomputer at Livermore will be more than 3 times as large as either of those machines. These three machines are funded by the Office of Defense Programs as part of the ASCI program. The ultimate goal of this program is to build a computer capable of 100 trillion operations per second that will be used to model and simulate nuclear weapons as part of DOE's Stockpile Stewardship Management Program.[4] Such a machine, planned for completion by fiscal year 2004, would equate to over 25 times the total capability of all of DOE's current supercomputers combined.

The "Option White" machine planned at Livermore is expected to cost about $85 million to complete. The cost of acquiring other ASCI program supercomputers and eventually the machine capable of 100 trillion operations per second is unknown at this time. The cost of operating supercomputers is also expected to increase. Livermore officials told us that the ASCI Option White machine will consume over 6 megawatts of electric power when complete. Six megawatts is enough electric power to supply about 5,000 homes for one year. Consequently, Livermore projects that by fiscal year 1999, its supercomputing costs will increase to about $46.5 million, up from about $35.2 million in fiscal year 1997.

[4]The goal of the stockpile stewardship management program is to ensure the safety and reliability of the nuclear weapons stockpile without actually testing nuclear weapons. A key part of this program is the ASCI program, which is aimed at providing advanced computers and software to accurately simulate nuclear weapons' performance.

DOE Has Unused Supercomputer Capacity That Can Be Shared

DOE is underutilizing its supercomputing resources and is missing opportunities to share them. Consequently, laboratory contractors may be acquiring additional costly supercomputers while DOE still has capacity available that could meet their needs. With respect to utilization, we found that DOE's laboratories are utilizing, on average, only about 59 percent of their available supercomputer capability.[1] The rates of utilization we observed ranged from 31 percent to 75 percent. DOE is missing opportunities to use its available capacity in part because it no longer emphasizes that opportunities for sharing should be considered. The only exception to this situation occurs at the National Energy Research Supercomputing Center at Lawrence Berkeley, which is set up to be a user facility. At this site, sharing with off-site users funded by DOE is substantial. However, there is little sharing overall. With about 41 percent of its existing capacity—almost 1.7 trillion operations per second—unused, DOE is missing opportunities to better share supercomputers among sites as an alternative to buying or leasing new machines.

DOE's Supercomputers Are Underutilized

DOE is underutilizing its existing supercomputers. According to utilization data we obtained from DOE's laboratories, the average utilization rate is about 59 percent overall, which is low in comparison with the higher rates (70 to 75 percent) reported at Lawrence Berkeley and other laboratories. With substantial additional capacity being added in fiscal years 1998 through 2000, overall utilization rates may decline even further. Although DOE expects usage to increase dramatically when the ASCI program is further developed, the extremely large size of the ASCI computers means that if even a small percentage of the capacity of those machines is available for sharing, they could potentially meet all of DOE's other supercomputing needs. However, DOE cannot say whether or how much unused capacity it has because it no longer monitors supercomputer workloads and utilization, even when a laboratory is seeking funding for a new machine. Thus, the Department lacks basic information on how effectively the machines are being used.

Utilization Rates at Some Laboratories Are Very Low

The utilization rates for supercomputers varied widely at the laboratories we visited, from about 31 percent to about 75 percent (see table 3.1). Because DOE does not maintain this information, we asked DOE laboratories to generate utilization data for each of their supercomputers from available site records. Data were available for 35 machines at seven

[1] Within the computer industry and DOE, there is no standard definition of utilization. For purposes of this report, unless otherwise noted, utilization refers to the percentage of a supercomputer's available processors actually in use for a given period of time.

laboratories. While table 3.1 displays utilization for fiscal year 1997, data we obtained for fiscal years 1994 through 1997 showed similar results.

Table 3.1: Supercomputer Utilization Rates at Seven DOE Laboratories, Fiscal Year 1997

DOE laboratory (number of computers for which data were provided)	Average utilization in percent[a]
Argonne (1)	63
Berkeley (6)	75
Livermore (11)	51
Los Alamos (6)	39
Oak Ridge (4)	69
Pacific Northwest (1)[b]	31
Sandia (6)	57
Total (35)	**59**

Note: To obtain average utilization in percent, the number of processors in use at each laboratory was multiplied by the hours in use and then divided by the number of processors available multiplied by their total hours available for use. See ch. 1 for a more detailed discussion of how we measured supercomputer utilization.

[a]Los Alamos reported it did not have utilization data for two of its eight supercomputers. Sandia reported it did not have utilization data for 4 of its 10 supercomputers.

[b]Pacific Northwest National Laboratory, which was making the transition to a newer, larger machine, had the lowest utilization rate—31 percent—on its older machine. Data for the newer machine, which was still undergoing acceptance testing, were not generally available and consisted primarily of test runs and thus were excluded from this analysis.

Source: Seven national laboratories and DOE.

Utilization rates for individual machines varied because of a number of factors, including whether the machine is new or old, or an experimental or a more stable production model. Laboratories' supercomputer officials told us that in some instances, researchers prefer some supercomputers more than others because the preferred machines are more reliable or run the researchers' computer programs more efficiently. On the other hand, some machines at sites with particularly low utilization rates are old machines that are being phased out—such as the Pacific Northwest supercomputer listed in table 3.1—or new machines that are still being phased in. The highest overall utilization rate was at Lawrence Berkeley, which is a designated user facility available to any researcher in the United States funded by DOE's Office of Energy Research. This factor probably contributes to the relatively high utilization of the supercomputers at this facility and demonstrates the benefits of sharing supercomputing resources.

While recognizing that various factors can affect utilization rates, we nevertheless believe that utilization rates of 59 percent or lower show that these computers are being underutilized. Arguably, the threshold for underutilization could be set even higher, at 70 percent or more, since at least one site exceeds 75-percent utilization and several individual machines' rates exceeded 90 percent.

An alternative way of looking at utilization is to consider how many of the available processors are used to run very large jobs. The largest machines—more than 128 processors for this review—are also the most expensive supercomputers acquired by DOE. The acquisition of these machines is typically justified by the need for very large machines to run very large programs simultaneously across many processors in order to complete the work in a reasonable period of time. Ideally, most of these machines should be running very large programs most of the time or at least a significant percentage of the time available. If a facility does not have a significant number of large jobs to run, it may be more cost-effective to buy one or more smaller supercomputers to run the smaller programs and to look for the opportunity to share a large machine with another facility. In fact, the laboratories' data showed that the largest machines are severely underutilized. During 1997, less than 5 percent of the jobs run on the largest supercomputers at DOE laboratories used more than one-half of the available processors. In other words, these supercomputers are severely underutilized for the types of programs that were used to justify their acquisition.

In many cases, these larger machines are being used to run a large number of smaller programs that would have fit on smaller, less expensive supercomputers. In some cases, such as the ASCI computers at Livermore, Los Alamos, and Sandia, the computer programs needed to fully use the capability provided by these machines are still being developed. For example, through the end of November 1997, less than 1 percent of the programs run on Sandia's ASCI Red supercomputer, which can process 1.8 trillion operations per second, used more that one-half of the available processors.

Other laboratories may have more very large machines available than very large programs to fill them up. For example, Lawrence Berkeley has a large 512-processor supercomputer that is utilized 75 percent of the time it is available. However, less than one-half of 1 percent of the jobs run on that machine require more than one-half of its processors. Beginning in fiscal year 1999, Lawrence Berkeley plans to replace this 512-processor supercomputer, ranked as the fifth most powerful supercomputer in the world when acquired in fiscal year 1997 for $26 million, with a newer model estimated to cost $27 million and capable of up to 1 trillion operations per second. As previously discussed, the Lawrence Berkeley Laboratory is a designated user facility for all Office of Energy researchers and thus has a fairly high utilization rate. However, locating another large-scale supercomputer at Berkeley may call for careful evaluation, given that two other DOE facilities funded by the Office of Energy Research are already in the process of acquiring, or planning to acquire, new large-scale supercomputers in about the same time frame. As discussed earlier, Los Alamos is now acquiring and installing a $40 million machine with several thousand processors capable of 1 trillion operations per second. The same careful evaluation is called for in the case of the Pacific Northwest Laboratory, which plans to replace its new machine (a 512-processor machine accepted in 1998) with a larger machine in about the same timeframe. Given the large amount of unused capacity at DOE facilities, the new capacity being acquired, and the limited number of large-scale programs that require these very large machines, acquiring a new large machine for Lawrence Berkeley, Pacific Northwest, or any other laboratory may not be justified. Furthermore, as discussed earlier, if even a small percent of the capability of the very large-scale ASCI machines is available for sharing, this availability could meet all of DOE's other supercomputing needs.

DOE Does Not Monitor Utilization Rates

In the past, under a DOE order on computer management and acquisition (Order 1360.1b), cancelled in September 1995, DOE's operations offices and headquarters information technology managers were responsible for collecting and analyzing workload and other performance data. They used these data to help ensure that the Department's information technology resources—including supercomputers—were being used to their maximum effectiveness. The order specified that, commensurate with program requirements, computer managers analyze performance data to define workload trends and identify problems. These analyses were to help them to adjust workloads, maximize return on investments, and assist in projecting future workloads, among other things. Under the old order, DOE's operations offices were routinely involved in overseeing laboratories' management of computers, including supercomputers.

However, as further discussed in chapter 4, DOE canceled the order on computer management and acquisition, replacing it with a more general order on information technology management (O 200.1, September 30, 1996), which considerably reduced DOE's oversight over laboratories' information technology acquisitions. Under the new order, operations offices are no longer responsible for overseeing laboratories' computers (including research computers/supercomputers) and no longer collect workload and performance data on them. As a result, DOE cannot systematically monitor existing utilization rates before investing in additional supercomputers. When a laboratory is seeking funding for a new supercomputer, existing workload and utilization rates are not routinely calculated or factored into the decision-making.

Opportunities to Share Supercomputers Are Being Missed

Given the amount of existing unused capacity and planned growth in capacity, DOE is missing sharing opportunities because it does not emphasize to its program offices and laboratories that they should be looking for them. Only a limited amount of supercomputer sharing occurs at DOE's laboratories. Most sharing occurs at the Lawrence Berkeley's National Energy Research Supercomputing Center, which was specifically created as a user facility and is shared among DOE-funded users from across the country. The facility, funded by DOE's Office of Energy Research, has six supercomputers, associated data storage devices, and other related hardware. According to officials, the facility serves about 2,000 users at the Berkeley Laboratory, other national laboratories, universities, and industry across the country. Some sharing also takes place at other DOE laboratories. Lawrence Livermore reported that over 30 users from Los Alamos and Sandia currently use its ASCI Blue supercomputer, and a variety of users from Livermore and Los Alamos use the ASCI Red supercomputer at Sandia. However, the Los Alamos and Livermore laboratories are in the process of installing their own ASCI-funded supercomputers, which will each be capable of over 3 trillion operations per second; in contrast, Sandia's ASCI Red machine is capable of 1.8 trillion operations per second. In all likelihood, this huge increase in capacity at Los Alamos and Livermore will decrease the use of the Sandia machine.

The amount of existing unused supercomputer capacity in DOE's laboratories indicates that opportunities for sharing are being missed. For example, in a May 1995 report, DOE's inspector general criticized the Department for failing to consider alternatives to buying a $13 million machine at Pacific Northwest's Environmental Molecular Sciences Laboratory.[2] The report stated that three other sites "already had the

[2]"Audit of the Department of Energy's Environmental Molecular Sciences Laboratory," DOE/IG-0371, Apr. 7, 1995.

computer systems that could fulfill the needs of the new Research Laboratory." This is the same supercomputer system that Pacific Northwest had already started planning to replace, even before it had completed its final acceptance testing of the machine and placed it in service. In addition, planned new capacity to be added in the next year may compound the problem. As shown in table 3.2, over 40 percent of DOE's total existing supercomputing capacity of over 4 trillion operations per second is not being utilized, and additional capacity of 1 trillion operations per second is planned for delivery within the next year.

Table 3.2: Existing Used and Unused DOE Supercomputing Capacity, and Planned New Capacity

Billions of operations per second

Laboratory	Existing capacity[a]	Unused capacity[b]	Planned new capacity
Argonne	45.00	16.6	None planned
Berkeley	506.40	128.8	53
Livermore	350.30	172.9	780
Los Alamos	762.90	463.6[c]	200
Pacific Northwest	207.30	142.6[c]	None planned
Oak Ridge	206.60	64.7	None planned
Sandia	2170.25	895.5[c]	None planned
Total	**4148.75**	**1685.30**	**1033**

[a]Capacity measured as total peak performance in billions of operations per second.

[b]Unused capacity equals total capacity minus utilized capacity.

[c]The laboratory kept no utilization data on at least one of its supercomputers. To estimate unused capacity for the laboratory, we applied the weighted average utilization for the remainder of the laboratories' supercomputers to those machines for which data were not kept.

Source: Seven laboratories and DOE.

The lack of emphasis on sharing may be especially true at DOE's weapons laboratories, where ASCI machines with huge capacities are being built. According to officials of the weapons laboratories and of DOE's defense programs, they look for opportunities to share supercomputers within the ASCI program, but they believe that sharing among DOE's programs and laboratories is limited by various technical factors, including the state of communications links between them and problems with alternating between classified and unclassified computer operations. While these may be legitimate concerns at the level of 100 trillion operations per second, which is envisioned for the future and discussed later in this chapter, in our view they are not legitimate concerns at the current level of operations, as demonstrated by the experience at Lawrence Berkeley. Currently, four supercomputers at Sandia Laboratory have been using less than 40 percent of their available capacity. In addition, Los Alamos is building two supercomputers in the same room, which when complete will both likely rank among the top 5 to 10 supercomputers in the world and cost a total of about $174 million. One system, the ASCI Blue Mountain machine, is designed to achieve a speed of 3.1 trillion operations per second at a planned cost of $134.4 million. The other system, the Energy Research program's Nirvana Blue machine, is designed to achieve a speed of 1 trillion operations per second at a planned cost of about $40 million.

While this effort at Los Alamos might appear to be an example of supercomputer sharing across programs, in fact the two DOE program offices involved have no formal agreement to collaborate in building or using the two machines. However, they initially told us that the goal in building the two machines at the same laboratory was to achieve synergy in the development of numerical algorithms, hardware, and software. DOE officials initially spoke of connecting the two machines to achieve a peak performance of up to 4 trillion operations per second. However, a recent statement by the head of the ASCI program raises questions about this collaboration. He told us that he would like to see the Office of Energy Research remove the "Blue" designation from its machine to make it clear that this machine is not associated with the ASCI program.

DOE also is adding outside ASCI capacity while unused capacity exists at the weapons laboratories. Total existing unused capacity and planned added new capacity within the three ASCI program laboratories (Livermore, Los Alamos, and Sandia) is substantial. Despite this, in February 1998, DOE, through its Lawrence Livermore National Laboratory, leased additional computer capacity of about 200 billion operations per second from the Pittsburgh Supercomputer Center (formerly funded by the National Science Foundation) for 1 year at a cost of $4.5 million, in order to support the ASCI Strategic Alliances Program.[3] DOE's existing unutilized supercomputer capacity at the time of the new lease was more than 8 times the added capacity the Pittsburgh facility would supply. In addition, the planned new DOE capacity scheduled to come on line in fiscal year 1998 alone is 5 times greater than the amount of added capacity leased from the Pittsburgh facility. While, this decision may have been made in part because the program has not resolved how it is going to provide access to foreign nationals working at its university partners, there

[3]The National Science Foundation, in an effort to cut costs, sought to consolidate its supercomputing efforts and withdrew its support for the Pittsburgh Supercomputing Center. The Center, in an effort to remain viable, sought support elsewhere, including from DOE.

appears to be sufficient other capacity available in DOE to have met some or all of this need. According to DOE officials, they performed an informal analysis of the available supercomputing capacity within the laboratories, for which there is no documentation, before Livermore entered into this $4.5 million contract. The lack of documentation is not surprising because, as noted earlier, DOE does not require its laboratories to keep utilization data and the program offices that make most funding decisions do not routinely consider such information or the option of sharing existing resources.

According to DOE, if used to their full potential, the supercomputers of the future will process and generate more data than can be effectively handled by DOE's existing communications infrastructure and thus could hinder the ability to share supercomputers among sites. As discussed in chapter 2, the ASCI program's ultimate goal is to build a supercomputer capable of 100 trillion operations per second, or over 25 times the capability of all existing DOE supercomputers. Machines of this scale will generate enormous amounts of data and could potentially overwhelm DOE's communications infrastructure if not adequately planned for. For example, ASCI officials at Livermore estimate that the classified wide area network that handles their transmissions is currently 100 to 300 times too small to support their highest computing needs in the future. Research is under way as part of the ASCI program to address this issue.

Conclusions

DOE-funded supercomputers are underutilized in terms of both the percentage of time they are being used and the size of the programs being run on them. We believe that two factors contribute to the underutilization of DOE-funded supercomputers. First, DOE does not monitor its laboratories' supercomputer workloads and utilization and does not require that such information be considered when deciding to acquire new supercomputers. Second, DOE no longer requires the contractors and universities that operate its national laboratories, nor its program offices that provide the funding, to address opportunities for sharing supercomputers when justifying the need for new supercomputers. At a minimum, we would expect to find documentation of (1) workloads and utilization rates and (2) sharing opportunities within DOE's existing supercomputer portfolio when the acquisition of a new supercomputer is being contemplated. Without considering such information, decisions to acquire new supercomputers are, in essence, being made in a vacuum.

We believe that there are opportunities for DOE to rectify the low utilization rates for DOE-funded supercomputers by increasing the general sharing of supercomputers among sites and by concentrating the very large programs at one or more of the existing supercomputers, which also are underutilized in terms of running very large programs. Such action could lead to a rise in the overall utilization rate for supercomputers and could result in the more effective use of the largest machines to run the programs that were the basis for their acquisition in the first place. Taking advantage of these opportunities could obviate the need to acquire as many supercomputers or supercomputers of the size currently planned. We make recommendations in chapter 4 that will address this issue.

Agency Comments and Our Evaluation

In its comments, DOE stated that processor utilization is only one dimension of massively parallel computing systems and does not account for the other factors, such as memory size, memory bandwidth, and input/output bandwidth, that could render a supercomputer "fully saturated" at well under a 70-percent utilization rate. While, we agree that these and other factors would prevent DOE from achieving 100-percent utilization, we did not state that DOE should or even could achieve 100-percent utilization. Rather, we concluded that DOE was missing opportunities to improve its low overall utilization rate because it does not monitor utilization or require that opportunities to share supercomputers be considered before making decisions to buy supercomputers. We continue to believe that DOE can improve its utilization of supercomputer resources and achieve an overall utilization rate greater than its current 59 percent rate. DOE's National Energy Research Scientific Computing facility located at Lawrence Berkeley routinely achieves rates of over 70 percent on its massively parallel supercomputer. DOE argues that such machines are not similar to its ASCI supercomputers because they are stable "production" machines. However, DOE is using commercially available technology to build the large-scale ASCI supercomputers, which are in many ways similar to other supercomputers using the same technology. If DOE could improve its utilization rate by 10 to 15 percent overall, it could save tens of millions of dollars in acquisition costs for new supercomputers.

DOE also stated that the 5 percent of the jobs using over one-half of the processors on the ASCI Red supercomputer at Sandia account for 80 percent of the utilization of this machine. DOE also stated that our conclusion that 41 percent of its overall supercomputer capacity is available for sharing was erroneous because of the 80-percent utilization rate cited for the ASCI Red supercomputer. DOE therefore concluded that its supercomputers (1) do not have available capacity to share, (2) are being used for large-scale applications, and (3) have unused capacity that is actually close to zero. DOE also stated that the sharing of the ASCI program machines is very difficult because of national security concerns.

We disagree. The ASCI Red supercomputer is used only 43 percent of the total available time, including its use for all large-scale applications. The 80-percent utilization rate cited by DOE represents the portion of the 43 percent total use devoted to large-scale programs—in other words about 34 percent. Thus, a large proportion of this machine, up to 57 percent of total available time, is still available for use by others. With regard to the sharing of the ASCI machines, they were originally planned and are being installed to allow just this type of sharing. The three ASCI supercomputers are designed to have both classified and unclassified modules that can also, after following proper procedures, be linked together to run the largest programs. In fact, one of the requirements of the ASCI Red supercomputer was that it could be switched between classified and unclassified uses in less than 30 minutes. In addition, ASCI program documents state that 10 percent of the capacity of these machines will be available to users from outside DOE's laboratories, such as the universities participating in the ASCI program's research.

DOE's Proposed Changes May Not Improve Oversight of Supercomputers

DOE has not effectively overseen the acquisition and use of supercomputers, and its proposed implementation of the Clinger-Cohen Act will not improve its oversight. The Department does not have a process in place to ensure that supercomputer acquisitions are fully justified and represent the best use of funds among competing priorities. Instead, its existing program planning, project management, and budget formulation processes focus more on overall research projects than on the acquisition of supercomputers that support those projects. As a result, new systems are planned and acquired without DOE oversight, while substantial unused and underutilized capacity already exists within DOE.

In April 1998, DOE outlined plans for a new process to comply with the Clinger-Cohen Act, which requires that federal agencies implement a comprehensive, efficient approach to acquiring and managing information technology. DOE's new process separately manages administrative and scientific computers, leaving the responsibility for scientific computers—including supercomputers—to individual program offices. As envisioned, this approach may allow DOE's program offices to continue acquiring supercomputers outside the Department's normal process for implementing the Clinger-Cohen Act. This approach, contrary to what is envisioned in the Clinger-Cohen Act, effectively places the vast majority of DOE's information technology resources outside the purview of the Department's chief information officer.

DOE has established criteria for designating projects as "strategic systems" if they cost over $400 million, are an urgent national priority, are high risk, have international implications, or are 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 costs and schedules. The ASCI program will cost about $4 billion from fiscal years 1996 through 2010, is an urgent national priority because of national security concerns, and has international implications because it is a major factor in United States' support of the Comprehensive Test Ban Treaty. In addition, the program is high risk because it seeks to advance the state of the art in supercomputing and simulation well beyond current capabilities, has already experienced delays, has had its projected costs increased, and depends on as yet unknown technologies for success. However, the program has not been designated as a strategic system.

DOE's Existing Processes Have Not Effectively Overseen Supercomputer Acquisitions

Neither DOE's existing processes for research planning nor for overseeing information technology focus on the acquisition and use of supercomputers in an independent, comprehensive manner. Consequently, as discussed in earlier chapters, no one person or office within DOE knows how many supercomputers are at the national laboratories, what they cost, or how they are being used. As a result, new systems are planned and acquired without departmental oversight, while substantial unused and underutilized capacity already exists.

Existing Planning Processes Do Not Focus on Supercomputer Acquisitions

DOE's program offices, including its Office of Defense Programs and Office of Energy Research (the program offices that acquire most of DOE's supercomputers), conduct their own, largely independent, research planning efforts in keeping with their separate program missions. These offices have research planning activities that generally include the following similar steps:

- continuously redefining programmatic and mission needs,
- developing and submitting written research proposals (which may include a proposed supercomputer acquisition), and
- reviewing and selecting proposals for inclusion in DOE's program planning and budget formulation processes.

As these steps indicate, these processes focus more on overall research initiatives than on the specific supercomputer acquisitions that may be included in the initiatives. Furthermore, these activities are not standardized or systematically documented in either of the two program offices or in DOE as a whole. In practice, in the Offices of Defense Programs and Energy Research, research ideas develop in a variety of ways from different sources. Neither office has standardized procedures for reviewing and selecting proposals. Consequently, the Department does not have a systematic framework for weighing competing supercomputing proposals when they are included in research programs.

The results of the program offices' planning activities are to be integrated into the annual budget cycle. In this process, proposed research projects are included in "field work proposal packages" from each national laboratory and subjected to reviews by the operations, program, and budget offices; the chief financial officer, and the Office of the Secretary. Approved projects are incorporated into DOE's proposed budget, which is subject to review and approval by the Office of Management and Budget (OMB) and the Congress.

In DOE's process, proposed supercomputer acquisitions may not show up in budget documentation and thus are not systematically weighed against one another. For example, the acquisition of the $40 million Nirvana Blue supercomputer at Los Alamos has been included in two larger initiatives, the ASCI program and the Interagency Nondefense High Performance Computing and Communications Program. While those programs have been highlighted in the budget, specific funding and justification for the Nirvana Blue supercomputer has not been highlighted. In other cases, funding and justification may be only partially visible in budget documentation. This is true of the ASCI Blue Mountain supercomputer at Los Alamos. According to DOE and laboratory records, total funding for this machine for fiscal year 1999 is $38 million, but only $2.8 million is visible in budget documentation. Program officials said that the remaining $35.2 million for this machine came from elsewhere in the ASCI budget.

Little Departmental Oversight of Supercomputer Acquisitions

Under DOE's current order on information technology management, "Information Management Program" (Order O 200.1, Sept. 30, 1996), the Executive Committee on Information Management consists of senior program and staff officers and the chief information officer, who has a nonvoting role. The executive committe and chief information officer oversee major information technology investments, and the chief information officer has the specific responsibility of overseeing the Department's information technology process. Under the order, the executive committee and the chief information officer exercise no controls over supercomputer acquisitions, which are essentially managed and overseen by the program offices and the laboratory management and operating contractors. Thus, over 80 percent of the information management assets funded by DOE are outside of the Department's information management structure, including most systems (including supercomputers) at the national laboratories.

This situation contrasts with past departmental practices. Under a former order, which was canceled in September 1995 ("Acquisition and Management of Computing Resources," Order 1360.1b) the following requirements were in place:

- Laboratories planning to acquire supercomputers were required to submit detailed implementation plans justifying the acquisitions to DOE's headquarters program offices and the office of information resource management for review and approval;

- Laboratories annually submitted long-range site plans for information resources management to the program offices and the office of information resource management; and
- DOE's operations offices were required to determine whether laboratories, before acquiring additional supercomputers or other computers, were maximizing investments, taking into account use data on existing machines, and considering sharing computer assets.

According to DOE officials, the order was canceled as part of a departmental effort to streamline the management of the national laboratories and to eliminate unnecessary paperwork requirements.[1]

This lack of DOE oversight and controls over supercomputers means that even the most expensive systems are not necessarily visible to the Department's information technology managers. For example, DOE has not exercised systematic departmentwide oversight over five major planned or ongoing supercomputer acquisitions for Lawrence Berkeley, Livermore, and Los Alamos. These computers, funded by Defense Programs and Energy Research, have a projected total cost of well over $250 million for fiscal years 1998 through 2000. Similarly, we found that a planned $7 million upgrade of Pacific Northwest's supercomputer—proposed for fiscal year 1999 by the manager of Pacific Northwest's computing facility and included in a list of ongoing/planned acquisitions supplied to us by Energy Research—was otherwise undocumented within DOE's and the laboratory's ad hoc and formal planning processes.

DOE's Implementation of Clinger-Cohen Act May Not Improve Oversight of Supercomputers

The Clinger-Cohen Act requires that DOE and other federal agencies implement an effective process for investing in information technology. DOE recognizes that its existing procedures for acquiring information technology do not follow Clinger-Cohen criteria and decided in April 1998 to implement a new process for planning and overseeing investments in information technology. This new "dual track" process includes investments in both administrative and scientific information technology but subjects them to separate management. In so doing, the process recognizes that DOE's program offices have viewed supercomputers as research "tools," not as information technology. The new process is a

[1]In the past, the Office of Information Resource Management annually collected and reported data to OMB on major administrative computers and supercomputers being proposed for funding in a given fiscal year. However, under present OMB guidelines, the chief information officer no longer displays budget "crosscut" data on all major information technology—including supercomputers—in submissions to OMB with the annual departmental budget request. Instead, overall information technology costs in functional areas such as hardware and software are reported to OMB.

compromise. It attempts to implement the Clinger-Cohen Act but may allow program officials to keep their existing research planning processes and to continue to acquire supercomputers without subjecting them to any sort of overall investment strategy.

Clinger-Cohen Act Requirements

The Clinger-Cohen Act of 1996 provides criteria for federal agencies to follow when acquiring information technology, including supercomputers. Among other things,[2] the act requires agencies to implement a process for selecting, controlling, and evaluating information technology investments—a process that assesses and manages the risks of information technology investments on an ongoing basis. As part of the process, agencies are to develop and employ quantitative and qualitative criteria for comparing and setting priorities among alternative information technology investments. OMB guidance, known as the "Raines rules," lays out the investment criteria to be met.[3] The Clinger-Cohen Act also envisions a key role for the chief information officer, who under the act is responsible for, among other things, promoting the effective, efficient design and operation of all major information resources management processes for the agency; monitoring and evaluating the performance of the agency's information technology programs; and advising the head of the agency on whether to continue, modify, or terminate a program or project.

DOE's Proposed Implementation of the Clinger-Cohen Act May Not Follow Its Criteria

In April 1998, the Department decided to implement an investment planning and oversight process for major administrative and scientific information technologies. DOE's new process separates computers into two categories—administrative and scientific, which includes supercomputers—and establishes separate review and oversight processes for each category.

Under DOE's approach, proposals to acquire either administrative or scientific information technologies (above a threshold of $2 million per

[2]Under the act, the heads of executive agencies are authorized to procure information technology for their respective agencies (sec. 5124) and to provide and implement a process for maximizing the value and assessing and managing the risks of the information technology acquisitions of the agency (sec. 5122). They also must develop goals for the effective use of information technology and report to the Congress, as part of the budget submission, on progress toward those goals (sec. 5123).

[3]Criteria include having the proposed information technology system (1) support core/priority mission functions; (2) be undertaken because no alternative is as efficient; (3) maximize the use of commercial, off-the-shelf technology; (4) demonstrate a projected return on investment that is equal to or better than alternatives; and (5) employ an acquisition strategy that appropriately allocates risk between the government and the contractor.

machine) will undergo an annual review and selection process that DOE calls "dual track." As envisioned, projects will be screened by a steering committee (co-chaired by the chief financial officer and chief information officer, with program offices' resource managers as members), which would decide, in step with the budget cycle, which projects are to be reviewed on the administrative track, and which on the scientific (sometimes referred to as programmatic) track. Thereafter, the dual tracks are to be independent in the following way:

- For investments in administrative information technology, a project team develops a rigorous business case for the acquisition, obtains all stakeholders' input on requirements, and performs a cost-benefit analysis. Projects are then scored and ranked for technological risk, business benefits, and return on investment. Using this analysis, the Executive Committee on Information Management, acting as the corporate investment board, evaluates projects against broader executive priorities and makes selections for funding. During implementation, selected projects are to be monitored against performance measures established by the project team. In selected cases, post-implementation evaluations will also be conducted.
- For investments in scientific (programmatic) information technology, the process is less defined. According to the decision document for the process, these investments will not be evaluated using OMB's "Raines rules," but instead "program offices will plan and review these systems using appropriate criteria for research conducted by contractors." In addition, "the Secretary and OMB review [these] systems as part of the budget process." Also, under a new reporting requirement, scientific information technology is to be included along with administrative information technology in an annual report to OMB.

DOE's proposed process allows the program offices to retain their present processes for acquiring supercomputers and appears to categorically exempt supercomputers from DOE's normal process to meet the requirements of the Clinger-Cohen Act. According to a staff member in the office of the chief information officer, the precise details of the process for scientific information technology remain to be worked out among the program offices and the chief information officer. However, a Defense Programs official said that from that office's point of view, the agreed approach does not treat scientific computers as information technology nor subject them to any sort of oversight by the chief information officer.

It remains to be seen how DOE will implement in detail the Clinger-Cohen Act for supercomputers. On the one hand, the Department recognizes that its existing processes for scientific information technology may not follow the Clinger-Cohen Act's criteria. On the other hand, DOE's program offices view the act's oversight requirements as a potential impediment to their research efforts. According to program officials, supercomputers are basically research tools, not information technology investments. In addition, the program offices do not want the Department's chief information officer to play a greater oversight role over the supercomputer acquisition process, as envisioned in the Clinger-Cohen Act. They view the chief information officer as lacking in knowledge of their research missions. The newly approved departmental "dual track" process is a compromise by DOE to implement the act and yet keep scientific information technology (and supercomputers) in a special management category, not under the oversight of the chief information officer or the Executive Committee on Information Management. In this regard, the new process may allow the program offices to continue with their "old" supercomputer acquisition processes, which do not follow the act's requirements.

Another issue to consider in DOE's implementation of the Clinger-Cohen Act is that most (over 80 percent) of the Department's information technology funding is spent by its management and operating contractors that run most of DOE' major facilities, including the seven national laboratories. In this regard, the act defines information technology to include information technology equipment used directly by the agency and equipment used by a contractor under the following circumstances: The contract (1) requires the use of such technology or (2) requires the use, to a significant extent, of such equipment in the performance of a service or the furnishing of a product. However, the act also provides that the term information technology does not include any equipment acquired by a federal contractor that is incidental to a federal contract.

To date, DOE has not taken a position on whether it will argue that the Clinger-Cohen Act is or is not applicable to the Department's scientific information technology that is acquired and used by its management and operating contractors. However, according to DOE, the Department does not normally require its management and operating contractors to use a particular information technology in performing their contracts but leaves such matters to the contractors' discretion. Thus, according to DOE, scientific information technology, such as the supercomputers acquired and used by its management and operating contractors, arguably do not fall within the act's definition of information technology and are not covered by the act. DOE acknowledges that a narrow interpretation of the act's definition of information technology, even where technically and legally supportable, might not be well received by OMB and the Congress. DOE also recognizes that the argument that the technology is incidental is difficult to make when the contractors' expenditures related to information technology are high—as is the case at the national laboratories where DOE's supercomputers are located. Furthermore, in most cases, while DOE does not require the use of a particular system, it is clear from the nature of the work it is funding at the laboratories that they need supercomputers to complete the research. Thus, in our view, it would be inconsistent for the Department—given the size, cost, and importance to DOE's mission of the supercomputers, as well the laboratory contractors' expressed need for them to carry out their work—to argue that the supercomputers acquired by its contractors are not required to perform the contract or are incidental to the contract, and are therefore outside the scope of the act.

Some Supercomputer Acquisitions Need to Be Managed as Strategic Systems

DOE may not be managing its largest supercomputer acquisitions appropriately. DOE does not manage even the most expensive supercomputer acquisitions—such as the ASCI system—as strategic system acquisitions requiring the attention of departmental management at the highest levels. DOE has established criteria for designating projects as strategic systems if they cost over $400 million, are an urgent national priority, are high risk, have international implications, or are 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 costs or schedules. In prior years, the Department has not effectively managed such systems, which have often been late and over budget. DOE currently manages 11 projects as strategic systems, including two systems related to stockpile stewardship—the National Ignition Facility under construction at Lawrence Livermore (estimated to cost $1.1 billion), and the Tritium Supply Facility (total cost to be determined). No supercomputer acquisitions, including those for the ASCI program, are or have been designated as strategic systems.

Nevertheless, the ASCI effort to acquire a supercomputer capable of performing 100 trillion operations per second—to simulate the effects of aging and ensure the reliability of nuclear weapons—meets the criteria for being treated as a strategic system. The ASCI program is a separate line item in DOE's budget, will likely cost about $4 billion from fiscal years 1996 through 2010, is a key 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 supercomputing 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. Although these characteristics would appear to make the ASCI program a clear candidate for being designated as a strategic system, Defense Programs officials said they have not managed ASCI as a strategic system because it is a program, not a system, and does not meet OMB criteria for being treated as a capital investment in the budget. However, this position is not consistent with DOE's November 1995 "Joint Program Office Policy on Project Management," which noted that some strategic systems are actually programs that include projects. The ASCI program, which has at its heart an ambitious effort to acquire supercomputer systems, would qualify. The ASCI program has to date already spent, or committed to spend, $370 million on four supercomputers and will build two or more significantly larger ASCI supercomputers in the next few years. In terms of total program cost, systems acquisition cost, and other factors, the ASCI program appears to be a prime candidate for designation as a strategic system.

Conclusions

DOE has not exercised effective oversight of its supercomputers, and its proposed implementation of the Clinger-Cohen Act will not improve its oversight. Currently, no person or office within DOE knows at a given time how many supercomputers the national laboratories have, what they cost, or how they are being utilized. As a result, new systems are planned and acquired without departmental oversight, while substantial unused and underutilized capacity exists. This gap between capability and utilization may grow even wider as DOE acquires still more powerful and expensive systems. Consequently, DOE lacks assurance that its existing supercomputers are being efficiently and effectively used. The Department also lacks assurance that additions to this inventory represent a well-justified allocation of resources among the its competing priorities. Furthermore, DOE's proposed "dual track" process appears to categorically exempt supercomputer acquisitions from the Department's normal process for complying with the Clinger-Cohen Act. In addition, we believe it would be inconsistent for the Department—given the size, cost, and importance to DOE's mission, as well the laboratory contractors' expressed need for them to carry out their work—to argue that supercomputers are not required to perform a contract or are incidental to a contract and therefore are outside the scope of the act. Finally, the Department should keep in mind that its most important, valuable supercomputer systems need the oversight of top level management, whether as information technology investments, strategic systems, or both—simply as a good management practice.

Recommendations

Given the number and cost of DOE's existing supercomputers, the unused capacity that exists, and future planned acquisitions, it is increasingly important that DOE better manage the acquisition and use of these systems. Therefore, we recommend that the Secretary of Energy adopt an approach to information technology investment and oversight that meets the criteria set out in the Clinger-Cohen Act. Specifically, under such an approach, DOE should adopt a process for acquiring scientific information technology that (1) pertains to all Department-funded supercomputers; (2) ensures, prior to providing funds for the acquisition of any new supercomputers, that a written justification clearly demonstrates the need, addresses the benefits of acquiring the subject supercomputer, and allows for meaningful comparison with alternative investments; and (3) includes a laboratory-specific analysis of the utilization of existing supercomputers and an analysis of the potential to share supercomputers with other sites and/or programs. We further recommend that the Secretary designate the Department's most ambitious acquisitions of supercomputer systems—such as those in the ASCI program—as strategic systems warranting oversight at the highest departmental level.

Agency Comments and Our Evaluation

DOE disagreed with our recommendations. The Department believes that it has implemented an appropriate process for acquiring information technology—including supercomputers—that meets the intent of the Clinger-Cohen Act. The Department also believes it is unnecessary for ASCI to be designated a strategic system because effective program oversight is in place.

In its comments, DOE stated that it has taken steps to implement the Clinger-Cohen Act and has in place a comprehensive managerial review process for supercomputers. According to the Department, its implementation of the Clinger-Cohen Act recognizes that administrative and scientific information technology systems have different purposes and uses and therefore should be managed differently. Accordingly, under the dual-track approach, scientific systems such as supercomputers are to be reviewed by the program offices using appropriate criteria for research. DOE also stated that appropriate rationales and justifications for supercomputer acquisitions are developed during the annual departmental review of program budget proposals. As part of its Clinger-Cohen implementation, the Department will report "aggregate information" on major scientific systems through the chief information officer to OMB.

GAO agrees that administrative and scientific computers are used for different purposes. However, we do not agree that an appropriate Clinger-Cohen process for supercomputer acquisitions is yet in place or that supercomputer acquisitions by DOE's program offices should be exempt from departmentwide oversight. DOE's acquisition of supercomputers are not always visible in program planning or budget documentation, which tend to focus on the overall research process, not the acquisition of supercomputers, even if they cost tens of millions of dollars. In addition, DOE's efforts to implement the Clinger-Cohen Act through (1) its new dual-track approach for acquiring administrative and scientific information technology and (2) its plan to collect and report to OMB "aggregate information" on scientific information technology systems do not go far enough toward greater departmentwide oversight. In fact, the dual track approach supports the status quo by specifically excluding scientific information technology from oversight by the chief information officer and the Executive Committee on Information Management. This leaves supercomputer management to the separate program offices responsible for purchasing and using the supercomputers, which are not in a position to oversee and evaluate these systems as part of any sort of overall departmental investment strategy for information technology. Accordingly, we stand behind our recommendation that DOE should adopt a departmentwide process that meets the Clinger-Cohen Act criteria and includes supercomputers and other scientific computing resources.

In its comments, DOE also stated that our recommendation to designate the ASCI program as a strategic system was unnecessary, in part because the Department has a Clinger-Cohen type process in place. We disagree that an appropriate Clinger-Cohen process is in place, as discussed above. The process DOE is implementing in response to the act would allow the same program office that has a vested interest in acquiring a supercomputer to be the Department's oversight body for the acquisition of that supercomputer. In our view, this approach neither follows the act nor achieves the degree of high-level oversight that designation as a strategic system would provide. In this regard, considering that the ASCI program is critical to efforts to ensure the safety and reliability of the nation's stockpile of nuclear weapons, and meets all other criteria for designation as a strategic system, we continue to believe that greater oversight of the ASCI supercomputers is essential, whether in the form of (1) a comprehensive justification and acquisition process for ASCI and other supercomputers, (2) designation of the ASCI program as a strategic system, or (3) both.

Comments From the Department of Energy

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

The Under Secretary of Energy
Washington, DC 20585

June 29, 1998

Ms. Susan D. Kladiva
Associate Director, Energy,
 Resources, and Science Issues
Resources, Community, and Economic
 Development Division
U.S. General Accounting Office
Washington, D.C. 205048

Dear Ms. Kladiva:

Thank you for the opportunity to review and comment on the General Accounting Office (GAO) draft report entitled Information Technology: Department of Energy Does Not Effectively Manage Its Supercomputers. At the request of the Secretary Pena, the draft report has been thoroughly reviewed by the Department and the staff of the national laboratories. In brief, the Department believes the title and some of the content of the report are not warranted. The Department takes issue with two of the three findings and one of the two recommendations of the report.

See comment 1.

In the Department of Energy, supercomputing is a mission critical technology. The Department has driven the innovation that has secured the position of the United States as a world leader in enabling science through supercomputing. The Accelerated Strategic Computing Initiative (ASCI) is part of the science based stockpile stewardship program with a clear set of technical goals and milestones (see attachment). The ASCI technical strategy provides the necessary computing platforms and infrastructure to meet those goals. The metric for this program is not hundred percent utilization, but the benchmarking calculations on the path to supporting the comprehensive test ban treaty. In addition, in the Department scientific supercomputing requirements for various core missions are far different from administrative business computing. Statistics and facts related to supercomputing activities must be evaluated in this context to reach valid conclusions. To provide guidance for a more realistic assessment of the DOE supercomputing resources and their utilization, the three major **findings** that form the foundation of the draft report are discussed below.

See comment 2.

(1) DOE's Supercomputer Capabilities Have Grown Significantly.

The Department agrees that supercomputing capabilities have grown significantly to become a mission critical technology. Without the stated level of supercomputing capability, the Department would not be able to maintain the safety and reliability of the U.S. nuclear weapons under the stockpile stewardship program or meet other mission requirements.

(2) <u>DOE's Supercomputers are Underutilized and Opportunities to Share are Missed.</u>

<u>Underutilization</u>

"Overall, DOE's national laboratories used about 59 percent of their available supercomputing capacity in 1997. ...GAO found that during fiscal year 1997 less than 5 percent of the jobs run on the largest and fastest supercomputers at DOE's laboratories used more than one-half of the supercomputer's available processors."

In the previous generation of computing, machines became saturated (in the sense that no more users can be accommodated) in the range of 70% of the processor usage due to the complexity of staging multiple jobs and writing out large files among other things. In the era of massively parallel computing, our and other supercomputer users experience indicates that machines will become saturated at a substantially smaller percentage. However, processor utilization is only one dimension of massively parallel computing systems. In purchasing these systems, trade-offs are made to correctly balance processor performance, memory size, memory bandwidth, disk size, & I/O bandwidth. Choices are based on averages for a collection of applications. In any one application, some parts of the system will be fully used and other components less so to allow the swift, proper balance of data and information flow in the computer. For example, moving large data files initially into the computer's memories and later out of the computer's memories from disks and archival storage requires significant computer time during which the processors are mostly idle. This would be counted as unutilized computer time by the single, simplistic metric of available processor time used in this report and would be counted as time that is available for other users – both are inappropriate conclusions. Using the simplistic metric of available processor time, 100% utilization is impossible; in fact our utilization data and experience from the past 30 years indicates that utilization would be considerably less than 100%.

See comment 3.

As for the example cited where only 5% of the jobs used more than one-half of the supercomputer's available processors, in fact the 5% of the jobs accounted for over 80% of the utilization on Sandia's TeraOPS ASCI machine, which is the result intended. According to the Top 500 List the TeraOPS ASCI machine, which is manufactured by Intel Corporation, is the fastest and most capable machine in the world; it is capable of running important applications at speeds in excess of one trillion operations per second. This clearly demonstrates that the machines are dedicated to the type of job for which they were acquired. These facts demonstrate that the single metric of processor usage is inappropriate for determining system utilization in these advanced systems. Furthermore, these facts demonstrate that the mission systems are being used at utilization rates that are within expectations for leading edge supercomputing machines.

See comment 4.

<u>Missed Opportunities to Share</u>

"With the exception of the Lawrence Berkeley facility, there is only limited sharing of supercomputers among DOE laboratories and with DOE-funded offsite users.

> **Furthermore, at the end of 1997, about 41 percent of the DOE's total supercomputer capacity (equal to about 1.7 trillion operations per second) is unused and DOE plans to add capacity of another 1 trillion operations per second during 1998. Meanwhile, DOE is missing opportunities to share this unused capacity among laboratories because it does not emphasize looking for such opportunities."**

The conclusion that 41% of the supercomputer usage is available for other users or for sharing is erroneous. As the data above shows, the supercomputing systems are being used to accomplish large-scale mission calculations and therefore the true percentage of DOE's unused supercomputer capacity is close to zero. The 41% figure is based upon misapplication of a single utilization metric.

In addition, DOE runs a number of different types of computing facilities. The National Energy Research Supercomputing Center (NERSC), at LBNL, is a user facility comprised of production machines in the state of stable operations designed to accommodate a large number of users. These users are chosen specifically to explore the adaptation of specific applications to the experimental architectures of the machines. Centers at Oak Ridge National Laboratory (ORNL), Pacific Northwest National Laboratory (PNNL), and Argonne National Laboratory (ANL) are mission support facilities comprised of research machines designed for the solution of large complex scientific problems and to accommodate a small number of users. These machines are often one of a kind or first of a kind that are by their very nature less stable than production machines. For this class of machine, a large number of users and high utilization is not achievable. ASCI requires the use of machines even more advanced than these just described. Comparisons of the usage of these machines without considering the intended purpose of the machines lead to incorrect findings. Further, such sharing of ASCI machines is very difficult due to national security considerations.

See comment 5.

(3) <u>DOE Has Not Been Effectively Overseeing Supercomputer Acquisitions</u>.

> **"Overall, DOE lacks an investment strategy for acquiring supercomputers and does not follow the Clinger-Cohen Act criteria, which require that DOE and other federal agencies implement a comprehensive, efficient approach to acquiring and managing information technology."**

The Department's implementation of the Clinger-Cohen Act of 1996 was approved by the Executive Committee for Information Management (ECIM) during its meeting on April 23, 1998. The ECIM is the senior management forum for addressing Information Technology (IT) management issues in the Department. It is chaired by the Deputy Secretary and the membership consists of the heads of all Secretarial Offices. The Department recognized that there exists a distinction between investments in corporate administrative systems and investments made in programmatic scientific and technical systems. Corporate administrative systems provide data or services to a broad range of users and are almost always based on proven technology. However, scientific and technical systems are used in the conduct of research, which pushes the state-of-the-art, are typically used by a special category of user, and/or may be the object of high-risk research while they are also supporting the accomplishment of a mission. These differences require different management approaches for DOE corporate administrative systems and scientific and technical systems. For programmatic scientific and technical systems, program offices such as the Office of Energy Research and the Office of Defense Programs plan and review these systems using appropriate criteria for research.

The Secretary of Energy and the Office of Management and Budget review programmatic scientific and technical systems as part of the annual review of program budget proposals. Appropriate justifications and rationale are developed and presented during project and budget reviews. Information on scientific and technical systems with costs exceeding the ECIM threshold are reported through the CIO using aggregate information as specified in the proposed OMB form IT-1, accompanied by an annual statement by the cognizant secretarial officer that appropriate planning and review processes are in place and being followed.

Recommendations:

> **"Specifically, GAO is recommending that DOE adopt a Clinger-Cohen approach for acquiring supercomputers..."**

This recommendation was addressed above in finding number three.

See comment 6.

> **"...the Secretary designate the ASCI program as a "strategic system" warranting departmental oversight at the highest level."**

As to the recommendation that ASCI be designated a "strategic system", we believe that this change is unnecessary. The ASCI programs have extensive, well documented annual and long range plans that have continually involved senior management at the DOE, as well as computing and simulation experts throughout government, academia, and the private sector. The programs are performing well against these plans. We believe that the management systems are in place to assure effective program oversight together with informed, timely objective and well-documented decisions. Given the fact that we have the Clinger-Cohen process in place, this program receives high-level oversight it requires.

In summary, DOE has shown leadership in the United States in high performance computing. The machines discussed are cutting edge machines. Given the level of sophistication of the hardware and software systems, levels of utilization and current levels of sharing of these systems are adequate. DOE has in place a comprehensive secretarial review process for supercomputers that meets the intent of the Clinger-Cohen Act of 1996.

Sincerely,

Ernest J. Moniz

Attachment

The following are GAO's comments on the Department of Energy's letter dated June 29, 1998.

GAO's Comments

1. The letter included a colored attachment, which we did not include in this report.

2. In regard to DOE's comments on utilization, we agree that various factors would prevent the Department from achieving 100-percent utilization. However, we continue to believe that DOE can improve its utilization of supercomputer resources. DOE's National Energy Research Scientific Computing facility at Lawrence Berkeley routinely achieves rates of over 70 percent on its massively parallel supercomputer. DOE argues that these computers are not similar to its ASCI supercomputers because they are stable "production" computers. However, DOE is using commercially available technology to build the large-scale ASCI supercomputers, which are in many ways similar to other supercomputers that use the same technology. If DOE could improve its utilization rate by 10 to 15 percent overall, it could save tens of millions of dollars in new acquisition costs for supercomputers. DOE also asserts that it has utilization data for the past 30 years when, in fact, it stopped requiring the laboratories to keep such data in 1996, and no laboratory or DOE official made such an assertion or provided any such data during the course of our review.

2. While DOE points out that the 5 percent of the jobs using over one-half of the processors on its ASCI Red supercomputer at Sandia account for 80 percent of the utilization of this supercomputer, we note that the utilization rate for this supercomputer is only 43 percent. Stated another way, DOE is saying that 34 percent of the available time on its ASCI Red supercomputer is taken up by jobs using over one-half of the available processors. This still leaves significant unused capacity available to run other applications, including additional large programs. We therefore disagree with DOE's assertion that DOE's figures equate to "utilization rates that are within expectations for leading-edge supercomputing machines" and that the "true percentage of DOE's unused supercomputer capacity is close to zero."

3. We disagree with DOE's position on the percent of overall capacity available for sharing, and with the Department's view that sharing of ASCI supercomputers is difficult. In fact, as we point out above, the total use on the ASCI Red supercomputer, including the very large programs, is only 43 percent, and up to 57 percent is still available for other use. With regard

to the sharing of the ASCI supercomputers, they were originally planned and are being installed to allow just this type of sharing. To date, the three ASCI supercomputers are set up to have both classified and unclassified modules that can also, after following proper procedures, be linked together to run the largest programs. In fact, one of the requirements of the Sandia ASCI Red supercomputer was that it could be switched between classified and unclassified uses in less than 30 minutes. In addition, ASCI program documents state that 10 percent of the capacity of these supercomputers will be available to users from outside DOE's laboratories, such as the universities participating in the ASCI program's research.

4. In regard to DOE's comment on the implementation of the Clinger-Cohen Act, GAO agrees that administrative and scientific computers are used for different purposes. However, we do not agree that an appropriate Clinger-Cohen process for supercomputer acquisitions is yet in place or that supercomputer acquisitions by DOE's program offices should be exempt from departmentwide oversight. DOE's acquisition of supercomputers are not always visible in program planning or budget documentation which tend to focus on the overall research process rather than the acquisition of supercomputers, even those costing tens of millions of dollars. In addition, DOE's efforts to implement the Clinger-Cohen Act through (1) its new dual-track approach for acquiring administrative and scientific information technology and (2) its plan to collect and report to the Office of Management and Budget "aggregate information" on scientific information technology systems do not go far enough toward greater departmentwide oversight. In fact, the dual-track approach supports the status quo by specifically excluding scientific information technology from oversight by the chief information officer and the Executive Committee on Information Management. This leaves supercomputer management to the separate program offices responsible for purchasing and using the supercomputers, which are not in a position to oversee and evaluate these systems as part of any sort of overall departmental strategy for investing in information technology. Accordingly, we stand behind our recommendation that DOE should adopt a department-wide process that meets the Clinger-Cohen Act criteria and includes supercomputers and other scientific computing resources.

5. In its comments, DOE also stated that our recommendation to designate the ASCI program as a strategic system was unnecessary, in part because the Department has a Clinger-Cohen process in place. We disagree that an appropriate Clinger-Cohen process is in place, as discussed above. The process DOE is implementing in response to the act would allow the same

program office that has a vested interest in acquiring the supercomputer to be the Department's oversight body for the acquisition of that supercomputer. In our view, this approach neither follows the act nor achieves the degree of high-level oversight that designation as a strategic system would provide. In this regard, considering that the ASCI program is critical to the efforts to ensure the safety and reliability of the nation's stockpile of nuclear weapons and meets all other criteria for designation as a strategic system, we continue to believe that greater oversight of the ASCI program is essential, whether in the form of (1) a comprehensive justification and acquisition process for ASCI and other supercomputers, or (2) designation of the ASCI program as a strategic system, or (3) both.

Major Contributors to This Report

Resources, Community, and Economic Development Division

Ed Zadjura, Assistant Director
Dave Brack
Pat Dunphy
Dan Feehan
Jonathan N. Kusmik
Anne McCaffrey
Carol Herrnstadt Shulman

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

PRINTED ON RECYCLED PAPER

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