NA-ASC-113R-07-Vol. 1-Rev. 0

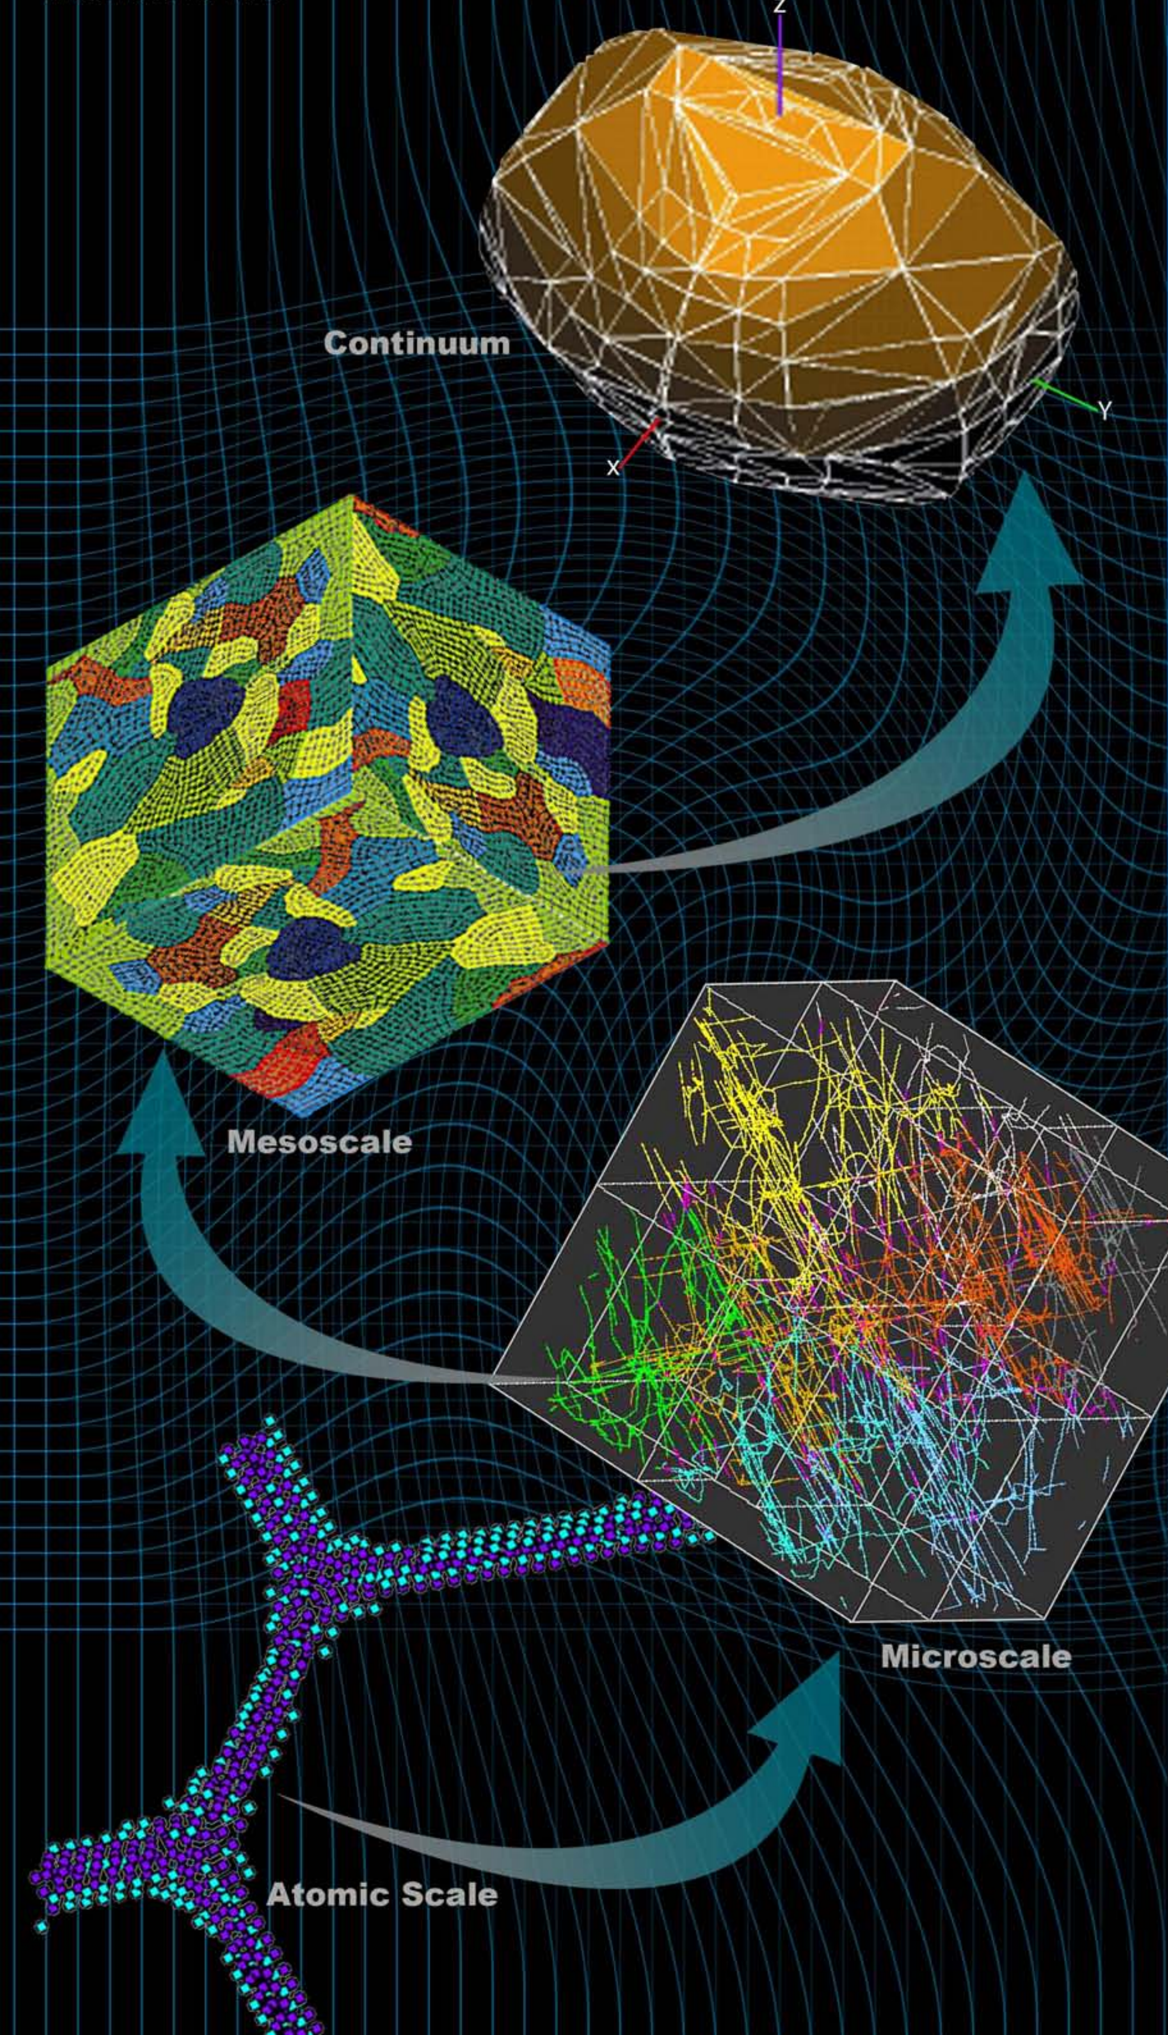

ASC™

A Platform Strategy for the Advanced Simulation and Computing Program

NNSA
National Nuclear Security Administration

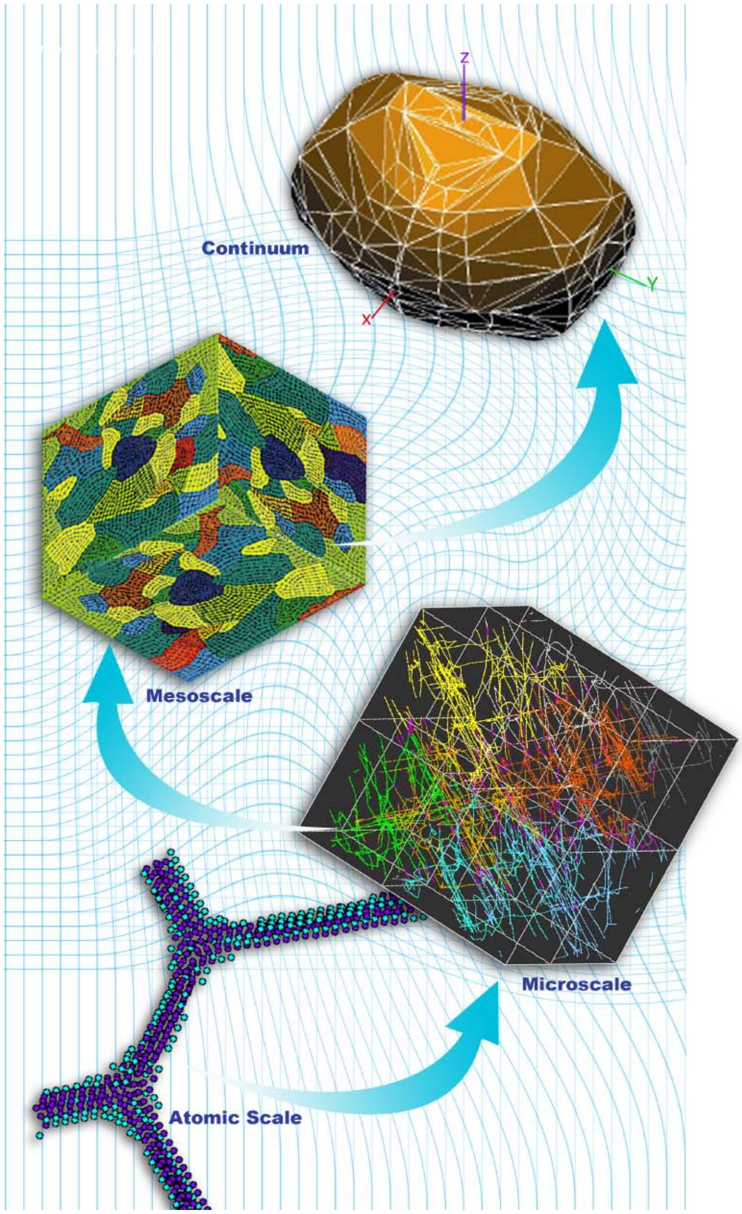

ON THE COVER:

High-performance computing platforms enable ASC to employ a hierarchy of models and modeling methods that range from the atomic to the continuum scale. Such high-resolution capabilities give us a better understanding of the physics and material processes relevant to stockpile stewardship.

Shown bottom-to-top are some examples from:

1. First principles and classical molecular dynamics investigating atomic defect mobility and interactions,
2. Microscale dislocation dynamics models studying single-crystal plasticity,
3. Mesoscale calculations for polycrystalline plasticity, and
4. Continuum structural mechanics examining properties of complex shapes. New ASC platforms will make possible capturing the response of materials and evolving structures over multiple time and space scales, a key step in establishing validated predictivity.

(Courtesy: Lawrence Livermore National Laboratory.)

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94-AL85000.

NNSA
National Nuclear Security Administration

SAND 2007-4343P
Issued by Sandia National Laboratories for NNSA's Office of Advanced Simulation & Computing, NA-114.

For more information, contact: Robert Meisner at bob.meisner@nnsa.doe.gov

A Platform Strategy for the Advanced Simulation and Computing Program

Author:

Robert Meisner, NA-114

Office of Advanced Simulation and Computing

August 2007

A Publication of the Office of Advanced Simulation & Computing,
NNSA Defense Programs

ACKNOWLEDGMENTS

I want to publicly acknowledge and thank those who made significant contributions to this document. Charlie Slocomb and Hans Ruppel were provocative sounding-boards for programmatic ideas and were instrumental in shepherding the document through the publication process. Major contributions to the overall technical content were made by Jim Ang (SNL), Ken Alvin (SNL on assignment to NA-114), Steve Poole (while at LANL), Adolfy Hoisie (LANL), and Mark Seager (LLNL). Jim Rathkopf (LLNL), Dale Nielsen (LLNL), and James Peery (while at LANL) made major contributions to the Stockpile Computing Requirements section. Steve Louis' (LLNL) and Mike McCoy's (LLNL) critiques added to the credibility of the final document. Reeta Garber (SNL) and her team, as always, provided excellent technical editing and an unerring artistic eye.

— *Robert Meisner*
NA-114

CONTENTS

The pictures at the bottom of pages 1 through 19 are a visual history of high-end computers sited at the NNSA nuclear weapons laboratories, starting from the present and going back to the very origins of computing. The laboratories have been instrumental in developing high-performance computing platforms and software and continue to be innovative drivers of supercomputing technology.

FOREWORD

The central problem that the ASC Program must address is replacing the need for underground nuclear testing with a science-based methodology based on predictive simulation validated by experiments and comparison to past tests. Understanding the myriad processes that make up a nuclear weapon explosion throughout the stockpile-to-target sequence (STS) is necessary to accurately predict detailed behaviors for any weapons system. Prediction requires extraordinary computational resources in terms of platforms, the supporting national infrastructure, and the experts to make it all work.

We have, in the ASC Program, developed competencies in secure high-end computing that have few equals in the world, and if we are to be successful in our mission, we must continue to stretch the boundaries of computing, always remembering that it is the national security mission that is, in the end, the reason for our vitality. While simulations are the end products that the Nuclear Weapons Complex employs to meet its national security mission, high-performance computing environments are the means to that end.

Today we are computing at a scale larger than we could have imagined only a few years ago. We can do science well at fractions of a petaFLOPS, on over 10^5 processors, with data streaming into the petabytes. In a few years, we will routinely use millions of processors on a single calculation, as multi-petaFLOPS systems find their way into the broader scientific community. PetaFLOPS today are essentially a matter of investment, and our hardware challenge is increasingly shifting to architectures in the exaFLOPS. With the large investments that we must make in computer systems — hardware, software, and infrastructure — comes the responsibility to ensure that we are using our resources wisely.

In this document, we discuss future directions for platforms and infrastructure that begin to address what ASC must become as a key part of the NNSA Complex 2030. This strategy is not a rigid contract that will remove our flexibility to solve problems, nor is it a way to park problems and ignore them until they become someone else's to solve. As we continue to build production-quality simulation capabilities, we are tackling problems concerned with cost effectiveness, resource sharing, system utilization, and future architectures and operating environments. It is our intention that this platform strategy forms the framework that will ensure our continued ability to address national security needs.

— Dr. Dimitri Kusnezov
Director, NA-114

EXECUTIVE SUMMARY

This *Platform Strategy* will guide future acquisitions of high-end computational platforms and their supporting infrastructure to ensure that we use our limited resources in a way that is responsive to the needs of the NNSA Nuclear Weapons Complex (NWC). It presents key ideas that guide our planning and that lead to cost-effective and useful implementations across the three defense laboratories that will be consistent with the *Complex* 2030 guidance. Six important points establish a solid framework for our future decisions.

1. Support NWC strategic plans:
   - Provide the responsive, efficient, integrated enterprise called for in Complex 2030;
   - Synchronize with the mileposts defined in the *ASC Roadmap*.

2. Define an enduring approach:
   - Acquire classes of computing—production and advanced—to support evolving stockpile requirements, not specific FLOPS goals;
   - Deploy strategy for capability platform design and siting that reduces the operating footprint within the NWC.

3. Align a balanced acquisition strategy with mission needs:
   - Ensure adequate computing to meet mission requirements by acquiring a cost-effective balance of capacity, capability, and advanced systems;
   - Team with industrial partners to ensure that performance of future computing systems keeps pace with increasing complexity of simulations.

4. Improve user productivity through prudent policies and new technologies:
   - Make capability and capacity production platform acquisitions available to users more rapidly;
   - Use consistent programming models supported by platform vendors to minimize the delay and expense of porting major ASC simulation tools to new platforms and minimize the interruption to users.

5. Increase Complex-wide integration and collaboration:
   - Create standards for hardware and software by executing tri-lab capacity procurements and building common software environments;
   - Operate production, capability platforms as National User Facilities that employ premium resources for NWC-wide priorities, not the host site's priorities.

6. Emphasize importance of a vital workforce and supporting infrastructure:
   - Recruit technical talent to keep NNSA at the forefront of simulation technologies;
   - Develop supporting infrastructure to ensure that platforms are optimally usable and productive.

1.0 INTRODUCTION

Recognizing the need in weapons codes to do large-scale 3D simulations, the Department of Energy (DOE) Accelerated Strategic Computing Initiative (ASCI) made a decade-long commitment to stand up a 100 teraFLOPS platform in 2004. Meeting this commitment was not the culmination, but the beginning of a new era in simulation. As 2005 dawned, the program now called Advanced Simulation and Computing (ASC) emerged with a vision of how predictive simulations would enhance our ability to meet current and future stockpile needs. ASCI had demonstrated the feasibility of using massively parallel computers as essential tools for exploring large-scale, integrated science questions and for answering questions about nuclear processes in weapons that previously could be addressed only through underground testing. Now, with the stockpile continuing to age, the questions become tougher, the demands on the codes more rigorous, and the computational requirements noticeably larger. While the Program name has changed, the essential nature of the computing challenge has not diminished.

ASC has transitioned from a program focused on feasibility to one focusing on mission-essential capability. With this transition, our goals have changed from procuring a specific level of technology (i.e., 100 teraFLOPS by 2004) to providing the agility to support current and projected stockpile requirements. Consequently, this document is not a plan to procure more supercomputers, but rather a platform strategy for meeting the needs of the stockpile. You will not find a plan to host a certain number of petaFLOPS at a given laboratory in 2010. However, you will find the need for a technical capability in a time frame that is linked to requirements identified in other programmatic plans and roadmaps.

The ASC Program depends on a well-balanced computing triad of applications, platforms, and supporting infrastructure to meet its nuclear security commitments. This document presents a strategy to ensure that we are using our resources most effectively and to guide future acquisitions of platforms and their infrastructure to be responsive to the needs of the Nuclear Weapons Complex (NWC). It presents key ideas to guide our planning that lead to cost-effective and useful implementations across the tri-lab community, implemented through the ASC Program Business Model, consistent with "Complex 2030" guidance and explicitly linked to the *ASC Roadmap*.

The National Nuclear Security Administration (NNSA) *Complex 2030* document[1] describes the "preferred planning scenario" that sets out the vision for the NWC of 2030. Central to this vision is a transformed NWC that is more responsive, efficient, and integrated. The nation will depend on the transformed NWC to develop and apply science-based stockpile stewardship tools to maintain deterrence without underground nuclear testing.

The *ASC Roadmap* provides details of the actual steps that must be taken to achieve significantly greater confidence in the calculations. In articulating the improvements needed in weapons science and science-

[1]*Complex 2030, A Preferred Infrastructure Planning Scenario for a Nuclear Weapons Complex Able to Meet the Threats of the 21st Century* September 2006, DOE/NA-0013.

based simulation to support transformation of the NWC, the *ASC Roadmap* calls out the essentials of mission-responsive computational environments and provides a guide for balancing platform, infrastructure, and science needs. Execution of the Platform Strategy is essential for the success of the Roadmap and, ultimately, stockpile stewardship.

The strategy presented here is strongly influenced by the history of high-performance computing (HPC) over the past decade both within the three defense laboratories and throughout the industry. Our experiences during the ASCI era, along with stockpile-related computing requirements that were documented through extensive studies done in fiscal year 2005, guide current and future platform configurations and policies. Our ability to acquire and effectively use the three classes of platforms described in this document will enable us to respond to the developing needs of the NWC, ensure that we have the resources to meet both the known and unforeseeable contingencies in the future stockpile, and provide the tools and infrastructure to meet future requirements of the Program.

1.1 Historical Perspective

Nuclear weapons tests were discontinued in 1992, leaving the U.S. without a direct means to certify the weapons in its stockpile. As a consequence, the Stockpile Stewardship Program (SSP) was created to maintain the safety, security, and reliability of the existing nuclear stockpile with improved experiments, surveillance, and advanced simulation capabilities. The advanced simulation component is being realized through the ASC Program, the successor to ASCI.

The ASCI code-development effort focused on the need for more predictive simulation through 3D representations of nuclear weapons phenomena, improved physics, and improved resolution. Legacy weapons codes were largely 2D and depended on the symmetry properties of the weapons designs coupled with an extensive underground testing program. The focus on development of 3D codes increased the need for more powerful computing systems with much more memory and many more processors. In fact, the architects of the early ASCI Program envisioned the need for computers that were more than 10,000 times as powerful as anything that existed in 1994.

The goals set for the Program during the initial planning activities of ASCI in 1994 were very aggressive, but they have all been achieved. In 1995, when the state-of-the-art of high-performance computing was measured in tens of gigaFLOPS (a gigaFLOPS is 10^9 FLoating-point Operations Per Second), we in the ASCI Program saw the need for hundreds of teraFLOPS (a teraFLOPS or TF is 10^{12} FLOPS) to begin to meet our long-term mission commitments. We recognized that to include the essential physics models, perform calculations in 3D, and ensure that the numerical approximations did not obscure issues associated with representation of the physics models, we would need 10,000 times the computing power that existed in 1995. The ASCI Program has, in the past decade, partnered with industry to acquire and deploy seven generations of supercomputers, from 1 TF (ASCI Red) to 3 TF (Red upgrade, Blue Mountain, Blue Pacific), 12 TF (ASCI White), 20 TF (Q), 40 TF (Red Storm, upgraded to 124.4 TF in 2006), 100 TF (Purple), and 360 TF (BlueGene/L).

LLNL
BlueGene/L/IBM
360 TF
2004

2.0 LESSONS LEARNED

The recent history of supercomputing at the national nuclear laboratories, characterized as the ASCI years, has resulted in valuable lessons that serve to guide the *Platform Strategy* as the Program proceeds along the path laid out in the *ASC Roadmap*. A critical subset of those lessons is explicitly incorporated into our current strategy.

Lesson 1: Productivity of the code developers, and designer and analyst users, is significantly increased if they are provided with a stable, modestly changing computing environment for applications development and production work.

Discussion: Our highest-end simulation needs have often required new technical features only provided by the most advanced vendor-integrated solutions—features often developed by partnering between the computer scientists and users at the defense laboratories and the computer vendors. To implement these solutions, we recognized the balance between short-term user productivity disruptions and the long-term productivity increases that result from more capable codes running on more powerful systems. Applying leading-edge technology to production use proved challenging and disrupted short-term productivity because of changes in operating environments and constant code reporting issues. To implement new technologies with minimum disruption to users, we developed a programming model that allows us to be independent of the characteristics of a particular generation of high-end computers. This model values optimizing long-term code portability over solutions that increase processor efficiency. Disruptions can also be minimized by deploying interim machines that let users address code porting, scaling, and tuning issues on a similar architecture and system software environment prior to the arrival of a leading-edge platform.

Lesson 2: The weapons system workload benefits from a mix of computer systems available to match cost-performance to problem needs.

Discussion: An initial focus of ASCI was the development and acquisition of high-end supercomputers to use on large 3D capability[2] problems. The experience of the ASCI Program with high-end parallel computing systems was that system usage was dominated by the users' need for immediate, short-term results that used a smaller number of processors and not by 3D, high-fidelity, full-system problems. This lesson was validated by workload characteristics established through several studies. Much of the workload comprises tens or hundreds of smaller and urgent, quick-turnaround jobs, which made it difficult to dedicate the capability computer systems to the largest problems and still meet commitments to deliverables.

Over the past several years, the marketplace evolved to a state in which high-end, commodity-based systems met the needs of many of our capacity[3] problems at substantially less cost. In the first 10 years of the Program, the focus was on commercial-off-the-shelf (COTS) processors, but the interprocessor commu-

[2]Capability computing is defined in section 5.1.

[3]Capacity computing is defined in section 5.2.

nications fabric and software were customized because there were no alternatives. Now COTS providers have expanded to include the communications fabric and systems software including the Linux operating system, as well as a variety of open-source software for debugging and performance tools, system monitoring and control, job scheduling, and file systems. This has made possible the acquisition of capacity computing systems, which can handle a substantial fraction of the workload at a much reduced cost.

Lesson 3: Investing in market-based supercomputers has proven to be a successful strategy for balancing system costs and progress in scientific computations.

Discussion: From its inception, the ASCI Program decided to work with the computing industry to leverage its business models to build supercomputers for scientific applications. In contrast to building one-of-a-kind custom architectures from which higher processor efficiencies might have been attained, the ASCI approach had several benefits. One was that commodity-based solutions provided an evolutionary path for applications, ensuring that code investments could cost-efficiently carry over to future generations. Another was that while market-based supercomputing platforms were expensive, they were still more affordable than custom-built architectures, and vendors were able to build, test, and deliver them in a relatively short time. Furthermore, given the low sales volume available from the scientific community, these business-based solutions leveraged a much larger market and provided a stable basis for producing ongoing generations of supercomputers.

While the ASC Program avoided investing in one-of-a-kind systems, it encouraged innovations in computer hardware and software that increased the capability and efficiency of high-end systems. Examples included participating in the design of the Red Storm architecture, contributing to the development of the BlueGene/L architecture, and advising on the design of future supercomputers resulting from the Defense Advanced Research Projects Agency's High Productivity Computing Systems project. These and other technical contributions to hardware and software systems components were integrated with the business plans of the vendors, increased the overall performance of the industry, and resulted in platforms that allow us to explore ever more challenging scientific questions.

Lesson 4: Bringing leading-edge systems to a production level is a time- and resource-consuming process that requires a strong partnership between the laboratories and vendors.

Discussion: To meet requirements, as well as to ensure that needed terascale computers would exist in the future, the ASCI Program procured systems that accelerated, and sometimes redirected, the business plans of its vendor partners. This resulted in both an invigorated HPC industry and a series of serial-number-1 systems acquired two to four years in advance of market offerings from a cross section of the industry. Over the past decade, the ASCI approach fostered competition and brought systems to market that would not have existed otherwise. Such systems provided a means to explore problem spaces previously not possible, but the application of such systems to production work introduced unforeseen problems in hardware and software reliability and system features.

Success in procuring leading-edge, serial-number-1 platforms and making them productive required healthy partnering with industry and significant expertise within the three laboratories to effectively deal

LANL
ASCI Q/HP
20 TF
2004

with adversity when it occurred during platform development and integration. The tri-lab integration teams worked closely with the vendor and the ASC applications groups to ensure that when the applications uncovered bugs in the hardware or software, the issues were dealt with quickly, and the solutions were implemented in a practical manner. Often understated, this was an essential ingredient for success in an advanced development environment. We found that developing long-term strategic partnerships with vendors gave us influence over system design and ensured that delivered systems were appropriate for the kinds of simulations required for the stewardship mission. These partnerships focused vendors on our application drivers and requirements and resulted in a more effective overall system.

Lesson 5: System software, essential to the overall success of high-end computer hardware, is particularly a challenge when we field unproven, serial-number-1 systems.

Discussion: In addition to partnering for system acquisitions, the ASCI Program fostered development of critical software component technologies through its PathForward industrial partnerships. The Program was successful in fostering development of system software components that were not part of the platform partners' business plans, but were essential to productive use of the systems and included scheduling software, debugging tools, system performance monitoring tools, and parallel file systems.

System software at the scale of ASC systems must run on hardware much larger than the vendor normally has available for code development and support. Special methods must be developed by the vendor to ensure that its software works on the delivered full system. This requires close collaboration between the vendor and the laboratory.

Lesson 6: Innovative architectural approaches provide significant future capabilities for the Program even though use of the advanced architecture may be confined, in the early stages, to a subset of the important physics simulations for which the Program is responsible.

Discussion: Advanced systems, though higher-risk endeavors, have provided significant returns to the Program in terms of our understanding of innovative architectural features. For example, the IBM BlueGene/L (BG/L) computer with 131K processors demonstrated efficient use of floor space and low power consumption, and the Cray Red Storm machine demonstrated high scalability with its advanced interconnect technology. Our use and investigation of the architectural innovations in Advanced Systems[4] with industrial partners ensures that we understand how to use the next-generation computers on stockpile stewardship problems and demonstrates to us that future HPC systems are suitable for our problems. The importance of these investigations to the Program has been supported by a National Research Council report[5] and an internal JASON study, titled "Requirements for ASCI."

For example, BG/L was initially focused on a class of science problems, largely molecular dynamics and turbulence modeling, both key to our understanding of basic weapons behavior. Recent integrated design

---

[4]Advanced Systems enable us to investigate new architectural features that show promise for achieving extreme speeds in addressing specific stockpile issues. Bringing advanced systems to market requires partnering with industry and sharing nonrecurring engineering costs prior to contracting for a machine to be built.

[5]*Getting up to Speed, The Future of Supercomputing*. National Research Council Report.

code results on BG/L show that future generations of BlueGene architecture, because of our technical input into the future designs, have the promise of widening their sphere of applicability and reaching the point of general-purpose machines to be used in full calculations with our modern codes. In addition, the RedStorm architecture has become a highly successful product line for Cray, which has already sold them to seventeen non-NNSA sites for science and national security applications.

We foresee that a model of strategic investments in advanced systems that begin life applied to a subset of stockpile issues but later expand in scope to become powerful general-purpose production engines is a viable model for ASC success. By fielding an advanced system (or an interim delivery of the next capability system), while still executing our day-to-day production computing responsibilities, we learn how to write code for such future machines, we begin to understand how our applications can be made to work efficiently on new (and possibly revolutionary) architectures, we begin to port the large codes in advance of the production phase, and we train users and system personnel in its use.

Lesson 7: Computing-at-a-distance is viable for doing even the most complex weapons system problems.

Discussion: Capability computers, whether located at Los Alamos, Livermore, or Sandia, have been considered to be tri-lab resources with major allocations to each of the laboratories. This is a successful model and enables scientists at each of the laboratories to effectively compute from their home laboratory on the most powerful systems available. This success has been enabled with sufficient classified networking resources and by data assessment tools that can be run in a variety of ways according to the particulars of the programmatic requirements and platforms.

Lesson 8: U.S. competitiveness and leadership in high-end computing, enabled by government investment and industry commitment, were necessary for progress in science-based stockpile stewardship.

Discussion: Science-based stockpile stewardship could not have succeeded without a sufficiently healthy high-end computing industry. Our particular mission for national nuclear security has required computing performance beyond the capabilities normally available in the commercial marketplace. Designing and delivering state-of-the-art supercomputers that met our stockpile stewardship needs within practical cost and schedule constraints were only possible through the efficiency and innovation of a healthy and competitive industry. Government support, through the ASCI Program and a limited number of other agencies, helped to engender this competitiveness and leadership through both competitive procurements and a commitment to share some nonrecurring engineering costs for development of capabilities needed to scale the systems to record-setting sizes. These high-visibility acquisitions grew out of a compelling mission need that the President and Congress supported through annual appropriations processes. U.S. companies in turn demonstrated their leadership through their senior managements' continued commitment to, and engagement with, the ASCI Program, despite the reality that, with some exceptions, ultra-scale HPC was not a significant revenue source for their companies. Even so, throughout this partnership, there has been a recognition within the industry that building supercomputers at ultra-scales does have a trickle-down effect on products that find wider use in other government and commercial applications, thereby providing competitive-edge tools that would likely not exist otherwise. The common denominator for continued U.S. competitiveness remains successful long-term partnerships between government agencies and industry.

LANL
Blue Mountain/SGI
3 TF
1998

3.0 STOCKPILE COMPUTING REQUIREMENTS

The capability to build and use large-scale simulation as a surrogate for nuclear weapons tests is a core competency essential for nuclear security mission success. Simulation is key for closing Significant Finding Investigations (SFIs), performing the annual assessment, assuring surety of weapons systems, certifying weapons modifications, designing the Reliable Replacement Warhead (RRW), being able to attribute responsibility for a nuclear event, and understanding weapons from rogue states. These simulations range from computations that require significant fractions of the largest computers in the world (thousands to tens of thousands of processors) to computations that are much smaller (tens to hundreds of processors). As the nation's nuclear stockpile moves further from the nuclear test base through aging, modifications made as part of Life Extension Programs (LEPs), or introduction of new designs such as RRW, the realism of ASC simulations must further increase through the development of improved physics models and methods requiring even greater computational resources. Stockpile stewards assess the uncertainties in their simulations of particular weapon systems by exploring multidimensional parameter spaces, which can require thousands of unique computations.

Problems at the high end of the computational spectrum have been a principal driver for the ASC Program and are moving us toward several petaFLOPS of computing capability by 2010 and an exaFLOPS computing capability by 2018. This is because of an as yet unmet need to significantly improve the overall predictive accuracy of our simulation tools. This need is becoming ever more urgent as we move further out in time from our last underground nuclear tests. Table 1 relates how the simulation focus and problem complexity have evolved since the inception of the program. The improvements in modeling fidelity are continuing to drive a rapid increase in problem complexity, which is more than doubling each year.

Table 1. The evolution of problem complexity in stockpile stewardship simulations

	Principal Simulation Focus	Growth in Problem Complexity
Early ASCI *1996–2000* *(4 yrs)*	Early demonstration of 3D simulations with coarse resolution	Increase dimensionality from 2D to 3D **Total computational increase:** **50X – 100X over 4 yrs**
Late ASCI – Early ASC *2000–2006* *(6 yrs)*	Increase 2D and 3D resolution, continue efforts to improve physics models	Decrease average mesh spacing by 2X – 10X, increase physics model complexity by 2X **Total computational increase:** **100X – 200X over 6 yrs**
Future ASC *2006–2020* *(14 yrs)*	Transition from calibrated to science-based physics and increase 2D and 3D resolution to meet predictive capability requirements	Decrease average mesh spacing by 4X – 8X; significantly increase physics model complexity (up to 100X) **Total computational increase:** **5,000X – 50,000X over 14 yrs**

LLNL
Blue Pacific/IBM
3 TF
1998

The overall plan for improved predictive simulation capabilities is articulated in the *ASC Roadmap*, which will be integrated into the broader Predictive Capability Framework for Defense Programs. This framework defines a timeline with improvements that address the key computational uncertainties and *ad hoc* models that limit more truly predictive applications of ASC simulation capabilities. Common to the development of improved predictive capability is increased spatial resolution of our computational models, as well as increased sophistication in the underlying physics and material models used in the simulation codes. This increased model fidelity will require substantial increases in computing capability, well beyond the current levels. Figure 1 highlights the computational resources needed for increasingly complex simulations and illustrates by example a subset of the stockpile stewardship requirements that we must address over the next decade. Advances in the fidelity of the physics and the accuracy of the numerical methods, and therefore our confidence in predictions of weapons system performance are dependent on the level of computing that can be brought to bear. Errors associated with numerical approximations may mask physical phenomena and shortcomings in physical models.

In 2005, NNSA acquired a 100-teraFLOPS-level system, the Purple machine, which provided an entry-level capability to address stockpile issues in three spatial dimensions. Before Purple, low-resolution 3D computations were run on much smaller systems but took over a year to complete. The 100-teraFLOPS-level system is capable of running low-resolution 3D simulations with turnaround times that are more nearly commensurate with the pressures imposed upon the weapons design staff and more responsive to Program needs.

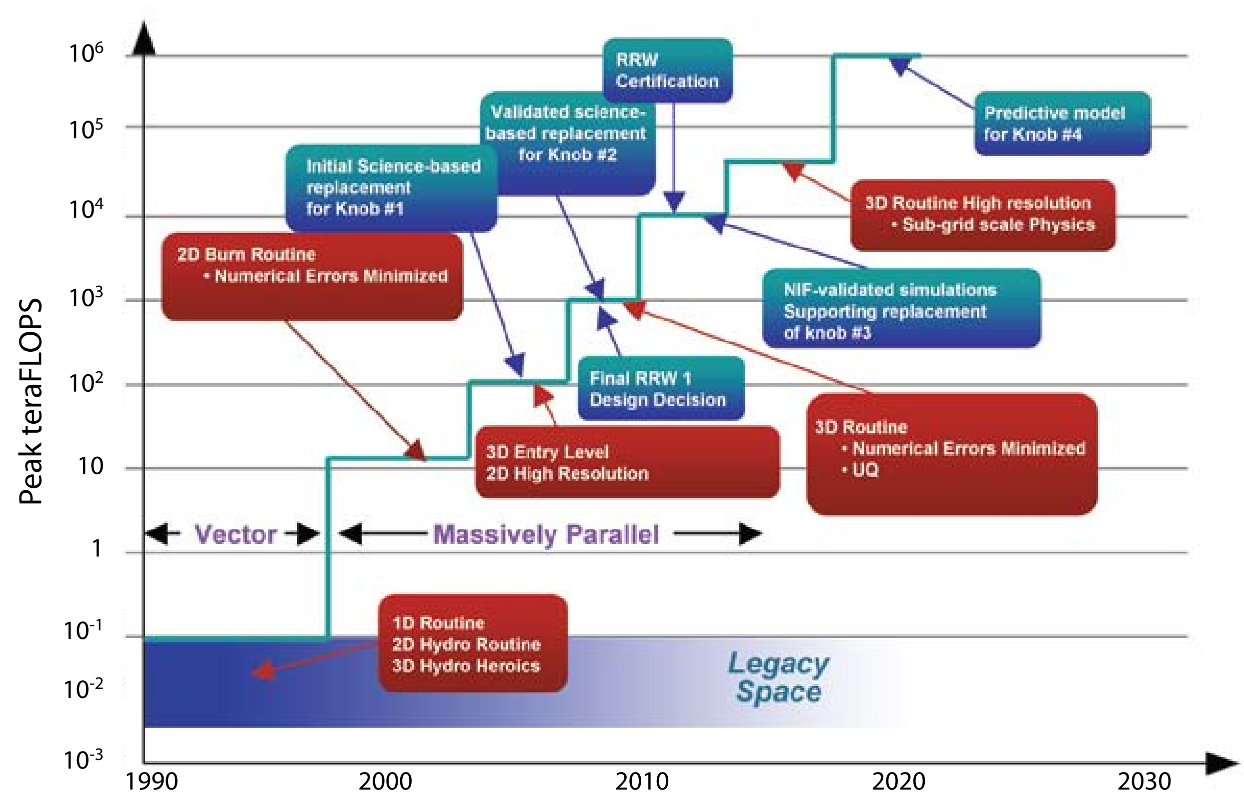

Figure 1. Computing requirements for selected weapons simulation characteristics

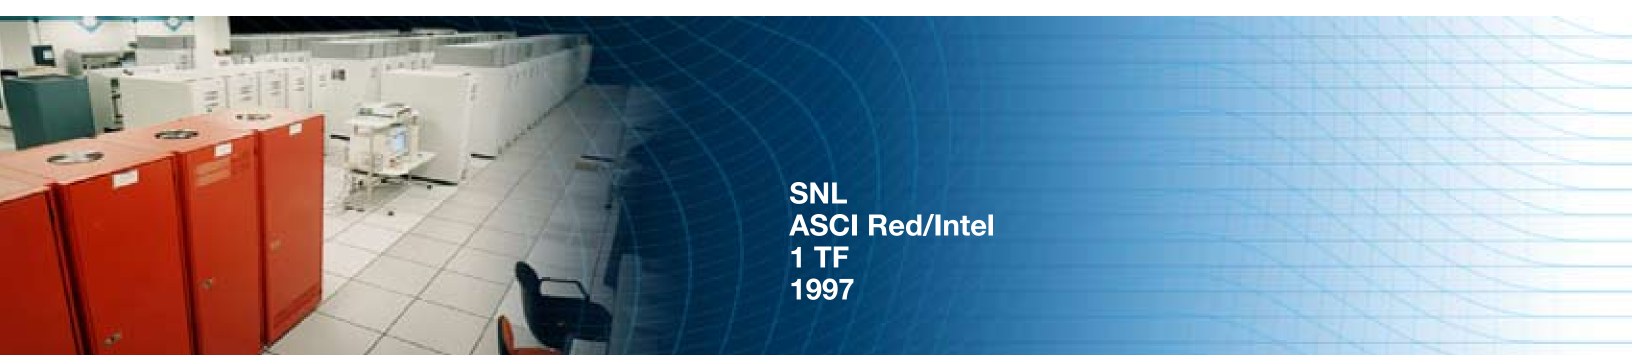

Recently, the first very-high-resolution 2D computations of critical phenomena have also been completed on the 100-teraFLOPS-level system. These results dramatically demonstrated that computational requirements for a single problem to achieve the necessary fidelity, even in 2D, can easily consume the largest computers available today. Repeating these simulations in 3D, which is necessary to eliminate the approximations inherent at the reduced dimensionality, would require orders of magnitude more computing capability. These simulations have shown that critical small-scale physics phenomena have now become accessible through high-resolution simulations, opening doors to scientific research and discovery that are essential to replacing *ad hoc* knobs[6] with science-based physical models. For example, at 100 teraFLOPS, we are able to replace the Knob 1 approximation with science-based, multiparameter representations of the device behavior.[7] This more realistic approach relies on a combination of enhanced scientific understanding and the computational power to represent it.

Two obvious methods of enhancing accuracy in predictions is by increasing spatial resolution, and, where appropriate, modeling in 3D rather than in 2D. The introduction of high-fidelity physics techniques also increases accuracy with accompanying increases in computational burden. Examples of applications of improved physics fidelity include the modeling of x-ray and particle transport, representation of the behavior of high explosives, and simulation of turbulence.

The modeling of the transport of both x-ray radiation and subatomic particles in the time evolution of a nuclear device is so compute intensive that crude, albeit computationally cheaper, approximations were historically used. A petaFLOPS platform allows the use of techniques that accurately represent transport phenomena of particles and radiation present in the high-temperature conditions of an exploding nuclear device and their interactions with matter in more nearly rigorous and mathematically defensible ways. Other basic phenomena, such as a representation of high explosive (HE) behavior through the solution of the coupled equations of reactive burn — the chemistry of HE behavior — also become possible at this level of computation power. A high-fidelity simulation of HE-driven implosion of a primary will be a large stride to better understanding the forces on the primary and the resultant nuclear burn phenomena.

Computational power at the 10 petaFLOPS level will enable us to use more accurate representations of turbulence, the microscopic behavior of materials, and detailed particle and radiation transport — all essential components of a confident representation of device performance. It will also enable representation of HE at the molecular and the grain scale and make possible the many thousands of high-fidelity simulations necessary to assess uncertainties through the exploration of multidimensional parameter spaces.

In the time frame of fielding of a 10-petaFLOPS-level machine (2011), NNSA may be certifying the first RRW weapons system. If this occurs, RRW will be the first system in 40 years to go into the stockpile without prior nuclear confidence tests. It will have components that use materials not previously built into weapons systems, and it will have extensive new safety and surety features. Certification will depend on the accuracy of simulations and the assessment of margins and uncertainties.

[6] Knob is an *ad hoc* model that compensates for gaps in our understanding of the physics of nuclear weapons systems.
[7] Science based replacement—2009 target for Focus Area 2, page 8, *ASC Roadmap*.

Both the *ASC Roadmap* and the *Predictive Capability Framework* emphasize delivering predictive simulation capability to the weapons complex by focusing on replacement of knobs. Use of knobs in code simulations can allow good agreement with experiment because the knobs themselves are calibrated to the experiments and/or nuclear weapons tests. Reliance on knobs must be eliminated as weapons deviate, either by design or natural aging processes, from the test data used for this calibration if we are to predict weapons performance without a return to testing.

The *ASC Roadmap* addresses the time frames in which we plan to replace the knobs with a combination of physics-based models, high-fidelity databases, advanced material constitutive property models, enhanced radiation transport capabilities, improved physical databases for relevant materials and processes, advanced opacity models together with other models as yet not realized. The replacement of the major *ad hoc* knobs is scheduled for 2009, 2012, 2015, and 2020.

These are aggressive goals that depend on the timely availability of the required experimental facilities, theoretical understanding, and computing capability to develop the underlying databases and to use them in sufficiently resolved 2D and 3D simulations.

Predicting the SSP's computing needs beyond the 10-petaFLOPS level relies on NNSA's experience with what has been required historically to make significant advances in reducing the uncertainty in our simulations. The requirements of expected future weapons systems for advances in predictivity achieved through advanced physics and material models, as well as increased resolution and routine 3D modeling, will likely drive our computing needs at least at the historic rate. For example, single, higher resolution 2D simulations now achievable on the newest capability platforms are requiring about one month of run-time at 100 teraFLOPS. Extending this resolution to 3D will require a 10,000X increase in floating point operations, or approximately 1 exaFLOP, with similar run times. Furthermore, when we factor in the need for additional resolution, or alternatively more advanced physics and material models, in order to replace calibration knobs with predictive capability, we can expect the computational needs to increase well beyond 1 exaFLOP.

In the early (ASCI) phase of the Program, the high end of computing capability, while well used on the smaller problems of the day, actually exceeded the needs of the complexity of codes and models, as developers and designers were learning to take advantage of new computational technologies. Currently, there is a good balance of computing and problem complexity, as code and model maturity have caught up with the capabilities provided as illustrated in Figure 2. This analysis compares problem complexity, as measured by computing resources required to run a single simulation, to available computing cycles on the high-end computer at that time. As this trend is projected into the future — based on estimates of the increased resolution and model complexity needed for the resolution of key physics knobs, and machines expected to be available at reasonable cost — we see that the growth in problem complexity will continue to challenge the capability of high-end computing to keep pace.

While this trend is of concern to the ASC Program, it should be understood as well that there are significant uncertainties in the projections. The current platform technology and investment strategy should allow the Program to roughly keep within reach of the exponential growth in problem complexity, subject to major changes in these trends.

SNL
Paragon/Intel
140 GF
1993

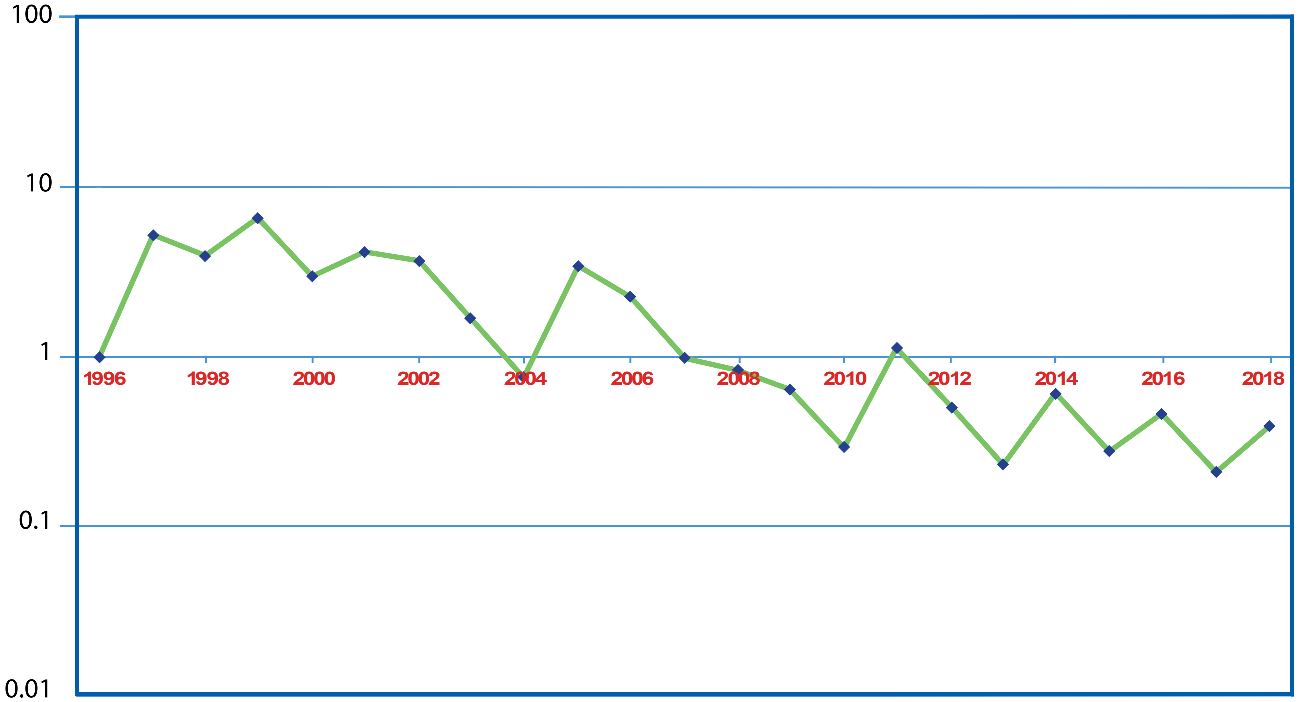

Figure 2. Ratio of computing capability to simulation complexity

Our ability to predict with confidence relies on the scientific research and discovery that reveal the physics and chemistry of weapons processes, physical models that capture processes, numerical methods that can faithfully solve the resulting mathematical equations, and computing hardware and software. Furthermore, productivity of our most valuable resource — designers, analysts, and code developers — requires that the computing infrastructure can turn around calculations within a time frame consistent with programmatic constraints. Responsible oversight of the national nuclear weapons stockpile demands a plan for the future that provides a solid foundation for building predictive capabilities and production environments to accomplish the work outlined in this section.

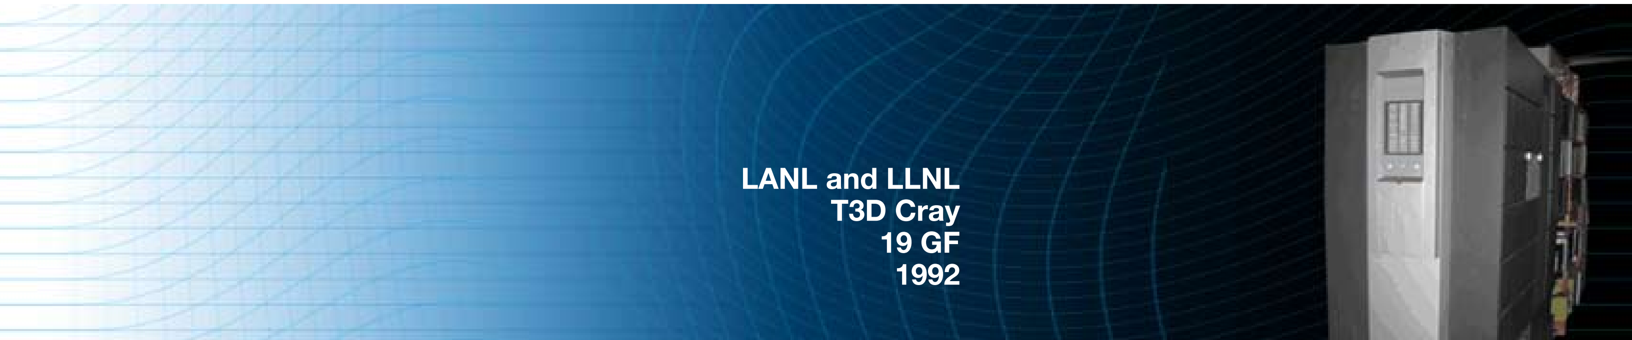

4.0 PLATFORM ACQUISITION PRINCIPLES

Computational platforms are an essential part of the tool set that ASC makes available to the weapons physics and engineering communities. While it is possible to run all problems on very expensive high-end computers, it is not the best use of resources, and large numbers of smaller jobs interfere with the timely execution of the largest problems because of oversubscription of computing systems. Four major principles guide our strategy for acquiring platforms to meet our mission needs as the Program balances needs and resources for solving today's problems while providing for more productive and cost-effective platforms for future problems. They are as follows:

1. **Maintain continuity of production**. It is necessary to maintain the productivity of the code developers and designers by providing architectures that, although they make increased computing power available, do not require all weapons work to slow while the codes are rewritten and ported to the new machines. This principle implies a conscious choice of continuity of infrastructure so that work can continue uninterrupted.
2. **Ensure that the needs of the current and future stockpile are met.** Two realities drive us to focus on the future while committing to get the job done in the present. One is that the complexity of the simulations we need to run is increasing as we transition from *ad hoc* model-based, calibrated codes to *ab initio*, physics-based codes. The other is that the supercomputing technology continues to evolve at a rapid pace. Together, these factors lead to the conclusion that future simulations are likely to be much different from those of today. Consequently, the ASC Program must strike a balance between making investments to meet current mission workloads and the imperative to be prepared for tomorrow's mission workloads.
3. **Balance investments in system cost-performance types with computational requirements.** Capability, capacity, and advanced systems offer a range of capabilities and costs to the Program. Simulations capitalize on the features offered by each at different costs. The ASC Program must invest in cost-efficient system types to match workload demands.
4. **Partner with industry to introduce new high-end technology constrained by life-cycle costs.** The Program must motivate industry to provide much increased capability and, as appropriate, drive the technology into new, promising, and applicable directions that have the potential to decrease time to solution and increase productivity by several orders of magnitude. However, the industry is now capable of building systems at scales that are beyond the reach of most operating budgets. With power consumption and footprint characteristics growing as dramatically as processor speed and system capabilities, the ASC Program will need to work with vendors to ensure that hardware designs take into account operating costs.

These principles guide our capital investments in production systems that maximize current productivity, and in advanced systems that are focused on future productivity improvements. Applying these principles moves acquisitions along parallel paths: we acquire incremental processing and memory improvements to our production capacity and capability platforms, and we work with industry to develop advanced systems with the necessary potential to improve productivity and/or reduce operating costs.

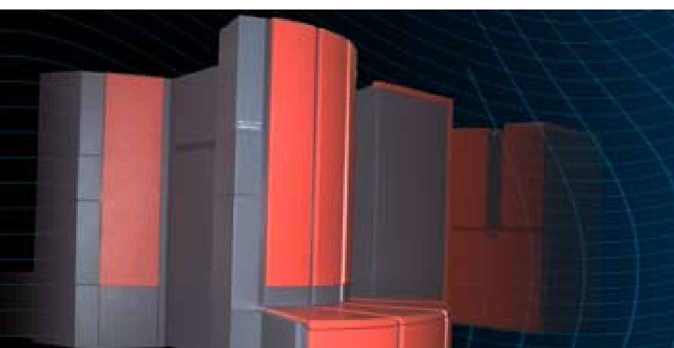

LLNL, LANL, SNL
Cray Y-MP
2.6 GF
1988

5.0 PLATFORM ACQUISITION PLAN

The previous sections outline the foundation upon which the ASC platform acquisition plan is built. The plan itself describes the approach the Program will use in providing future platforms for supporting the high-end simulation needs of the NWC. The approach is agile to respond to changes in the nation's nuclear triad and to the NWC of tomorrow, both as it is being defined by Complex 2030 and as it will naturally evolve based on the needs of the nation. It recognizes that current mission needs require that the Program continue to pursue supercomputing solutions that advance the state of the art for the nation.

This acquisition plan differs from the earlier ASCI approach, which demonstrated that massively parallel computing could be successfully employed for multidisciplinary, multiscale science, and which was based on a plan to invest in peak FLOPS on an accelerated schedule. The ASCI approach defined specific teraFLOPS goals to be achieved over a decade. The approach articulated here recognizes market directions (see Appendix 2, Technology Forecast) and intends to match the mission needs, with production platforms projected to be available. Furthermore, while this new approach takes advantage of the ability of the marketplace to meet many of the ASC Program's needs, high-end requirements still drive the need to partner with industry to develop novel solutions.

In this context, the acquisition plan projects that, with continued government investment, industry should be expected to provide general-purpose parallel systems with peak FLOPS characteristics as shown in Figure 3. This figure provides perspective for understanding the relative size of systems that we expect will be available to meet programmatic needs. It should not be interpreted as a program goal for procuring such systems. Further, it is not meant to suggest that such systems will be available without government investments. On the contrary, it depends critically on sustained investment by NNSA and other agencies with a vested interest in HPC.

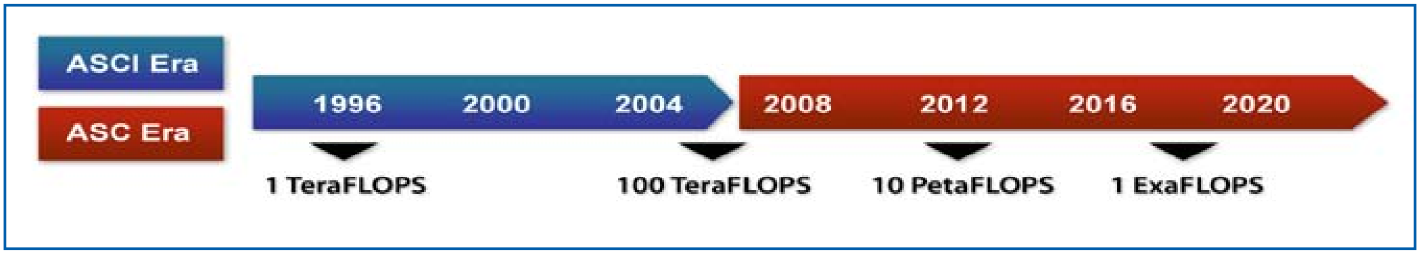

Figure 3. Time frame for availability of Peak Floating Point Operations per Second supercomputers

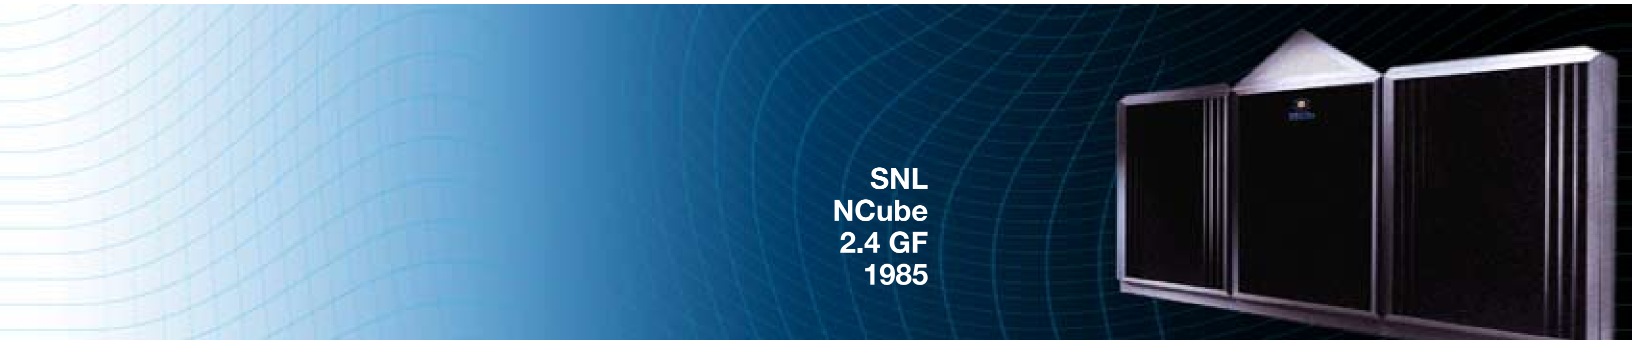

In FY06, the ASC Program adopted the Department of Energy's project management structure for acquisition of capital assets[8] for capability-class and advanced systems. An essential step in this process is defining and gaining approval for the mission need for platform procurements. A collateral process, the Future Years Nuclear Security Plan (FYNSP), identifies funding needs for the Program over five years. Planned platform procurements are projected in the context of the ASC Program's Roadmap and FYNSP funding profile. Projected deliveries over the life of the current FYNSP are shown in Appendix 3, Projected FYNSP Acquisition Plan.

The Program will use a mix of different systems to achieve its mission, that is, rely on relatively low-risk capacity and capability systems in stable production environments and use advanced systems as enablers for new scientific explorations and as pathways to increased productivity. The degree to which advanced systems are expected to support calculations will depend on many factors inherent in the acquisition approach, including architectural novelty, market availability of components, and maturity of programming models and application codes. During any single year, ASC will invest in all three types of computing to various degrees. The amount of investment in each type will vary, depending on available resources and mission needs. The essential characteristics of the three major classes of ASC computing are described in sections 5.1, 5.2, and 5.3 and summarized in Appendix 1, System Characteristics Summary. While it is convenient to think of Program acquisitions as falling into a single computing category, procurements may exhibit traits of more than one class.

Performance modeling and synthetic workloads both play important roles in ASC computer acquisitions by analyzing the performance of the delivered and implemented systems and through matching the mission needs to technology offerings. We will continue to use these methods to develop requirements for new hardware and software acquisitions as well as to evaluate the vendors' response to requests for proposals. Our objective is to ensure that a vendor proposal is suitable and architecturally well balanced for the ASC workload. After delivery, performance modeling and synthetic workloads will be used to check that the system is performing according to expectations, to identify problems that are reducing performance, and to tune key algorithms. Modeling, in particular, is a powerful tool for supporting objective decision making by allowing accurate, quantitative performance predictions based on a projected workload at every step in this process.

This strategy does not specify where platforms will be located, but all ASC resources will be viewed as national resources. Siting will be implemented through a careful and considered process. Many elements go into informing these decisions, including mission needs, available infrastructure, cost effectiveness of operations, and continuity of production computing. Siting will be determined during the decision process to initiate an acquisition. Furthermore, with the ASC Program transitioning away from independent computer center operations, the lead design team laboratory may not be the host for the resulting system.

[8]DOE Order 413.3A, Program and Project Management for the Acquisition of Capital Assets, 07/28/2006.

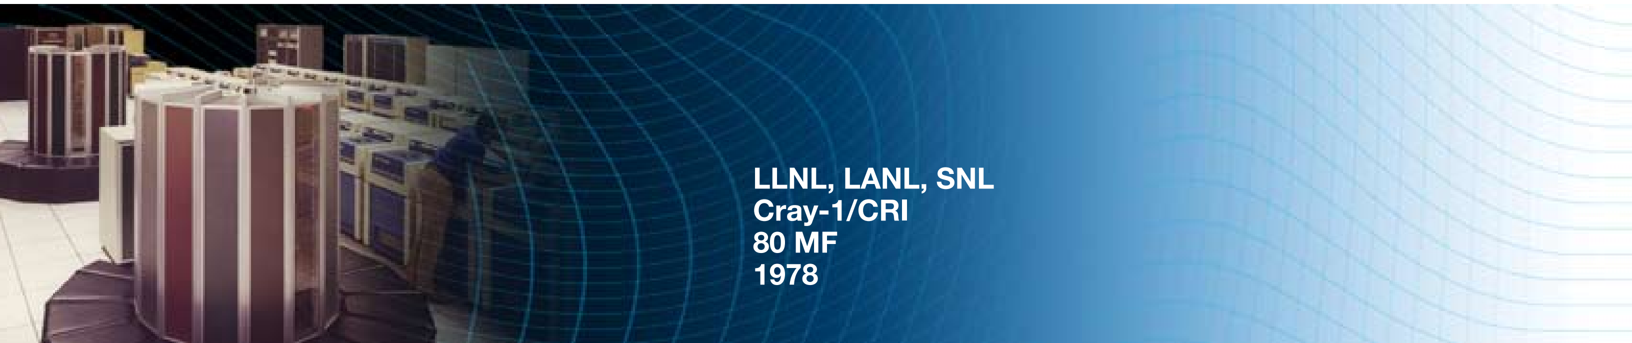

Treating HPC platforms as national resources dictates that their time will be allocated according to the mission needs of the NWC, not the host site. Systems designated to handle the production capability workload of the NWC will be prescribed by the ASC Platform Governance Model (see Appendix 4) as they transition to general availability. Time on capacity systems may be allocated to sites other than the host site if special needs are present, but sharing of capacity systems across the NWC will not be the usual case. Advanced architectures may be used by all the laboratories to get experience on novel architectures and to participate in their evaluation of value for ASC needs.

The useful life cycle of these systems depends on several elements: time-to-solution of current problem mean-time-between-failure (MTBF) of the system cost to maintain and ability to upgrade. The time-to-solution of a particular problem set on an older computer may be too long compared to the required turnaround because of processor limitations, memory limitations, or interconnect speed. In addition, the cost to keep the system on-line increases: as the system ages, the hardware maintenance issues become more onerous; vendors are more reluctant to keep the software features up to date; and the electrical energy required to effect a solution becomes much more expensive than for more modern systems. After a period that varies depending on the category of computer system, it is usually more cost effective and user productive to replace the computer with a modern system.

5.1 Capability Investments

Capability systems will be acquired primarily as general-purpose production systems dedicated to the most challenging problems of the NWC. This requires that the capability system be a leadership-class machine—that is, it will be among the largest systems in the world at a given time with the computational power, memory size, and interconnect speed necessary to solve complicated weapons system problems. The special challenges of capability systems mean that we may be willing to make strategic investments in features that will further the state of the art. Necessary features in capability systems that are difficult to implement because of the large scale of these systems include multiple cores; very large interconnection networks; reliability, availability, and serviceability (RAS) features to help with diagnosing problems; operating system capabilities to support performance and reliability at scale; and scalable input/output (I/O).

Capability computing systems are managed as tri-lab resources similar in value and uniqueness to large experimental facilities. Major programmatic computing efforts for capability systems are organized as computing work packages and are reviewed and prioritized for relevance, importance, and technical merit by a Capability Planning Advisory Committee (CPAC). Each proposed work package, called a Capability Computing Campaign (CCC), consists of at least one major calculation needing a significant proportion of an ASC capability system, together with related supporting jobs of smaller sizes. The cost of these systems is such that there may only be one platform procured for capability-class work in production in the NWC at any one time. This limitation is to control the size of investment in capability computers within the NWC so that the mortgages in the out years are aligned with available resources.

LLNL, LANL, SNL
CDC-7600
10 MF
1969

We will select capability systems that cause minimum disruption to users and that can be made available to a subset of users as soon as possible after delivery followed by the full community within a year of initial integration. Over the next 10 years, we expect to see capability systems migrate away from custom vendor software environments and more toward open source, which should further reduce the impact of long integration times on the user community.

The life cycle for capability systems is about five years.

5.2 Capacity Investments

Capacity computing is accomplished through the use of smaller and less expensive high-performance production computing systems that run parallel problems with more modest computational requirements. Capacity systems are typically low-risk technology and are the primary work tool of designers who now routinely work at 1,000+ processors. The requirement for capacity computing is largely driven by Directed Stockpile Work (DSW) — including model and code development, LEP, and SFI — code validation, and physics and engineering design and analysis. It is necessary that the Program make a sufficient investment in capacity computing to ensure cost effectiveness of the overall computational resource to minimize the technical risk to the Program and to ensure that capability platforms can be largely dedicated to computing at significant scale. We will determine the size of the capacity computing resource through accurate tracking of current usage and pent-up demand.

Capacity computing systems are now based on commodity Linux clusters that are in use throughout the scientific community. Capitalizing on the commodity nature of capacity systems, ASC has developed a procurement strategy that builds a common hardware environment across the three defense laboratories over multiple years. The objective is to quickly build, field, and integrate many Linux clusters of various sizes into classified production service through a concept of Scalable Units (SU). This approach should dramatically reduce the overall Total Cost of Ownership (TCO) of these systems relative to the best practices in Linux cluster deployments today.

To further reduce ASC Program costs, we are standardizing on a common Linux computing environment as a base. This common environment will help reduce redundant efforts across the laboratories and will make the systems available to users more rapidly, thus increasing the lifetime value of these systems. Since capacity computer systems are mostly comprised of off-the-shelf components, they should generally be made available to users in less than four months. One major advantage of having a tri-lab capacity computer procurement is that the systems can be acquired less expensively and can be available more readily to the user community. Further, with a consistent user environment and common operating system for capacity computers across the NWC, it should be possible for users to do their capacity work at any of the three sites that have available resources.

Our experience is that the life cycle of capacity systems is about four years.

LLNL, LANL, SNL
CDC-6600
3 MF/s
1964

5.3 Advanced Systems Investments

Advanced systems extend the limits of technology by exploring promising architectural approaches. A major element of the advanced systems element of our strategy is to understand future systems that will likely be quite different in nature from current capability and capacity systems. Acquisition and use of these machines enable the Program to examine new directions that have promise for future general-purpose weapons system computing, while at the same time applying new technology to current problems. These advanced systems may require new ways of approaching our problems and may require significant changes to weapons codes. The pay-off is the promise that we can move past "Moore's law" restrictions that are looming and threatening to curtail the speed increases that we have seen over the past decades and that there will be follow-on systems that can use the code modifications that we have made to address our future stockpile simulation obligations. This advanced systems strategy helps lead us far into the future, limits our vulnerability to technological surprise, and helps both the NWC and U.S. industry maintain a differentiating competitive advantage. It also attracts well-qualified people to work in the Program with the promise of being able to make new contributions at the frontiers of supercomputing.

While advanced systems are relieved of the "time to general availability" requirements of production systems, they are still expected to support a portion of the NWC workload. The design and building of advanced systems are higher-risk endeavors that push vendor partners to tailor their business plans for significant potential rewards. As such, advanced systems require partnering with industry over an extended period of 5 to 10 years. These partnerships typically will have an applied research and engineering phase that will result in early technology demonstration systems prior to final capability deliveries. The collaborative partnership also could result in a longer productive lifetime by designing the advanced system architecture to accommodate a mid-life upgrade option. Acquisitions will usually be structured to account for evaluations of technology readiness and early deliveries, and will contain sufficient options for either partner to terminate the relationship. The desire is that advanced systems projects progress to final delivery stages, although delivery of final capability is an option predicated on technology's ability to meet the contracted needs of the SSP.

We will, where it is appropriate and when it meets the needs of the ASC applications, collaborate with other government agencies in the acquisition process to leverage limited funding resources (cost sharing both between agencies and with industry) and HPC expertise. While a healthy industry depends on competing ideas and investments, individual agencies do not generally have sufficient resources to single-handedly execute advanced systems projects. Where application spaces overlap, opportunities are created to build advanced systems that otherwise would not be possible. In addition, a wide intellectual audience produces more capable products. As a result, products come to market that strengthen U.S. competitiveness in HPC and the science it enables. While the ASC Program may pursue advanced systems in support of the NNSA mission, our intent is to collaborate with and leverage the larger HPC community.

LLNL, LANL, SNL
IBM 7094
120 KF
1960

6.0 SUPPORTING INFRASTRUCTURE AND PEOPLE

Although ASC computing systems are at the heart of the computing capability provided to users, the integration of these systems with their supporting infrastructure and the people who do the integration and make the overall system operate efficiently for the user community play an equally essential role in effecting the platform strategy.

Balanced with the production computing machines, there must be sufficiently large storage and a networking fabric that permits rapid movement of data, both as simulation output to storage and directed to users' desks; there must be powerful and user-friendly visualization and data analysis capabilities to render and distill the massive amounts of data into useful forms; and we must provide the ability to access computers at a distance. Both production and advanced systems are tri-lab resources and must be available to all the sites for weapons production workload and for code development and porting.

Throughout the history of large platform acquisition and deployment, we have found that each new generation of architecture increasingly challenges the expertise of both systems programmers and modern code developers. Now, and certainly in future generations, we will need far greater knowledge, experience, discipline, and creativity on the part of those charged to make the high-end systems usable, efficient, and extensible. The need for increased competence goes hand in hand with increased complexity of systems. People with the level of computing expertise that we need in ASC are a resource that must be continually built upon to ensure the future health of the Program.

An important relationship between ASC and the open scientific computing world has been the Academic Strategic Alliance Program, now called the Predictive Science Academic Alliance Program. This program, spanning five universities over the last 10 years, exposes young graduate students and postdocs to the power of interdisciplinary, large-scale scientific computation. ASC not only funds research in disciplines relevant to the weapons program, but also makes available open computing resources to students and their mentors involved in these activities. We have found this to be a source of technical fertilization to our verification and validation efforts and for new ideas in computer science and algorithms. It is also an important aspect of recruiting future generations of young engineers into national security programs.

The workforce at the defense laboratories is crucial to the success of the ASC Program. We must build on our alliances and additional networks to other academic and industrial communities. We must recruit the staff that will keep our Program at the forefront of HPC and enable us to provide a simulation environment that meets our nuclear weapons responsibilities.

LLNL, LANL, SNL
IBM 704, 705
5 KF
1955

7.0 SUMMARY

The platforms component of the ASC Program is faced with making choices among competing priorities and must select from available options with mission goals in mind. The constraints that the Program faces are limited resources to expend on the tri-lab computing infrastructure, minimizing the disruptive element of new architectures that require rewriting codes, and not unnecessarily imposing new programming models that create serious difficulties for code development. The overriding objective is to maximize the productivity of users and developers while at the same time providing the capability to enhance confidence in simulations of device performance outside the data-range provided by the nuclear test base.

The ASC Program, like its ASCI predecessor, recognizes a national responsibility to ensure that the commercial computing sector, for whom this is a small component of their business, continues to pursue technology advances that enable large-scale scientific explorations for both weapons and nonweapons related problems. The national security enterprise understands the need to drive the industry in directions to ensure the specific program-driven resources will be available when needed, and to influence, to the extent possible, new technology directions.

We see a technology watershed occurring in the availability of COTS assets, both hardware and software, which dovetails with the primary needs of the Program. Concurrent with major 3D calculations that may consume our largest capability system, many simulations essential for our stockpile deliverables are comprised of smaller runs that can be performed on less expensive, easier to field capacity systems. The avenues that industry is offering now provide opportunities to field an infrastructure that meets many user needs and that can be done within resource limitations.

We have benefitted significantly from the lessons learned in the first decade of the Program. We have drawn upon those experiences, folded in a new predictive simulation strategy, linked acquisitions to the objectives of the *ASC Roadmap*, accepted the resource constraints within which the Program must operate, and developed an acquisition strategy that encompasses our principles and an understanding of the projected technical environment. In summary, we have crafted a platform strategy to build the computational infrastructure to meet the needs of the nation's SSP through an appropriate mix of capability, capacity, and advanced systems.

LANL
Eniac
Twenty 10-Digit Decimal Numbers
1945

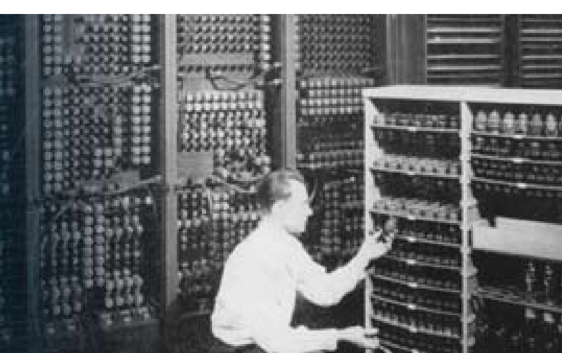

APPENDICES

APPENDIX 1.
System Characteristics Summary

The following table summarizes the salient characteristics of the ASC computing system categories.

	Function	Capacity	Capability	Advanced
Time to General Availability	Completed integration and made available to end computing users	Less than 4 months as a general rule	Less than 1 year	Often greater than one year
Reliability, Availability, and Serviceability	Built into hardware and software	Minimal features	Advanced	Advanced to minimal, depending on the architecture
$/FLOPS	Measure of proximity to commodity products	Low	High	Lowest to highest
Watts/FLOPS	Measure of electrical power required for computer	Medium	High	Low to high, depending on architecture
TriPoD	Common Linux computing environment	Yes	Possibly	No
Open Source		Extensive	Some	Some
Code Applicability	The degree to which existing codes are able to run on new platforms	Extensive	Extensive	May be limited (in early phases)
Period of Procurement	Length of time to procure one or more systems	2 year multiple systems	1 year for single system delivery	5–7 years single system with early deliveries to demonstrate concept
Life Cycle	Time to obsolescence of a system	4 years	5 years	4–6 years, depending on upgrade options
Global Memory Size		Small	Large	Small to medium
Risk	Risk to the program that a computer system will not be suitable for ASC codes without substantial changes to the programming model.	Low	Medium	High
Influence on Design	ASC ability to specify design specifications	Component-level	Node-level	Chip-level
Vendor Availability	Number of vendors who can provide the system type	Significant	Limited (3 or 4 vendors)	Limited (about 3 or 4)

APPENDIX 2.
Technology Forecast

This section of the ASC platform strategy includes a list of ten key technologies that are likely to impact the path to petascale in the 2010 time frame and exascale computing before 2020 on the ASC platform roadmap.

1) Over the time frame of interest for petascale systems, Moore's law will continue to govern. Moore's law encapsulates the empirical observation that the number of transistors on a die at constant cost doubles about every 24 months. However, Moore's law is silent on performance of those transistors. In the past, shrinking transistor feature sizes permitted a drop in circuit voltage and enabled an increase to the core's (CPU's) frequency at a rate that led to application performance doubling about every 18 months. However, complementary metal-oxide semiconductors (CMOS) feature sizes are now so small that continued frequency increases also seriously increase power consumption. On the other hand, using the smaller feature size to add more cores has little effect on power consumption. Thus, the majority of microprocessor performance boost in this time frame will come mainly from a geometric increase in cores and at a diminished rate from core frequency increases. Specifically, this means that processors will increase from dual core today to have 32, 128, or even more cores. The development of next-generation system architectures and a bridge to the ASC investment in distributed memory (MPI) applications must acknowledge and address this forecast.
2) Microprocessors with multiple hundreds of cores (CPUs) will force the microprocessor designers and system architects to address the balance factors for processor access to local memory hierarchies. With this many cores (CPUs) per processor die, it will be very easy for a large fraction of the total computational capability of the processor complex to lie idle while waiting for data, and that will force an effort to address the memory wall again. Understanding the balance between the need for larger and more capable memories and increased computational speed will be imperative for ASC codes to perform at an optimal level. As we increase physical fidelity and detail in our codes, improve numerical algorithms, and increase resolution, the need for more memory (as well as more capable memory) will continue to grow. We need to understand better (perhaps through algorithmic modeling) what demands our applications continue to make on memory subsystems and what the balance between processing and memory should be. Having application performance data, both simulated and measured, will allow us to give factual input to memory and system vendors. A critical effort will involve the development of new heterogeneous (implicit inner-loop/explicit outer-loop) parallel programming models and support for them.
3) From the scalable system architecture perspective, support within processors for addressing distributed versus local memory can address critical bottlenecks that limit scalability for integration into petascale systems. Because changes and new functional capabilities in the memory management units (or directly in the memories) will be required for multi-core processors, ASC and the greater HPC community may have an opening to provide suggestions for memory subsystem functionality that will be useful for the integration of multi-core processors into scalable systems. Having processor memory subsystem controllers that are designed to support concurrency will be a tremendous capability and increase the base level performance of the ASC codes and the broader HPC code base as well.
4) Scalable system architectures will also drive requirements for improvements in the interconnect fabric to take advantage of potential memory subsystem improvements. Specific areas for improvement include the interconnect bandwidth, latency, and message injection rate. In addition, this time frame

will likely see the development and use of optical technologies for the interconnect fabric. Future scalable system interconnect bandwidth requirements will accelerate the practical viability (i.e., price) of optics for interconnect technologies. Additional drivers include the weight of copper cables and the imposition of shorter distance limits as signaling rates increase on copper cables.

5) A new development that could have a significant impact on the ability to achieve exascale levels of performance before 2020 is the incorporation of Single-Instruction-Multiple-Data (SIMD) or vector processing accelerators on scalable system compute nodes. Achieving scalable performance on these heterogeneous architectures will be possible only if the coupling between cores and accelerators is extremely tight. It will also require a continuing focus of interdisciplinary efforts to develop a new generation of parallel algorithms and their associated advanced solvers that are able to circumvent the interconnect and memory subsystem bottlenecks between the compute nodes and their integrated accelerators. The recent announcement of the ASC/Los Alamos National Laboratory collaboration with IBM is an example of current attempts to attach accelerators (cell processors) onto nodes of large-scale systems.

6) Early in the Program, ASC standardized on a programming model for distributed memory with explicit MPI message passing. The investment that ASC made in application software and associated algorithms and solvers that use this programming model paid dividends in the ability to port ASC applications across most of our capability and capacity systems with relatively modest levels of effort. As the industry moves to parallel applications at unprecedented levels, the Program will need to explore the ability for departures from our current programming model to address performance and scalability issues on advanced systems. The ASC Program may need to make a decision to change its applications portfolio to the next programming model to allow applications to take full advantage of the computing potential offered by an advanced architecture. This decision will probably be driven by the need to improve the parallel efficiency. When the difference in parallel efficiency is great enough, we will have a technical rationale for this shift. This decision has already been made in other HPC markets and will be monitored.

7) Scalable system software is a critical enabling technology for future petascale systems. We expect to have full-service operating system software like Linux and lightweight kernel (LWK) operating system software. ASC invested in two systems that use LWK system software, Red Storm/XT3, and BG/L, which have demonstrated scalable and reliable performance up to the full scale of their respective systems. While the pursuit of application performance at the petascale and beyond is likely to continue to require the use of an LWK, there are users that want to use at least some functionality provided with a heavy-weight operating system. The DOE/ASCR's FAST-OS Program is supporting a broad range of approaches to span the gap between full-featured Linux operating systems and stripped-down LWK operating systems and runtime system software. The move towards a large number of cores may also support the development of new ways to distribute the workload that supports the notion of flexible specification of services to balance scalability/performance with services/productivity. There are also several LWK-based operating systems being developed in the open-source community that might be of value.

8) Scalable, parallel file system technologies are critical enablers for scalable systems and also as the integrating element within a simulation environment consisting of simulation engines (capacity and capability systems), data assessment engines (data manipulation and visualization systems), and archival systems. Looking forward, the way to achieve I/O performance targets for petascale systems may be through larger aggregations of devices and links. A stumbling block to such levels of performance is the number of devices that will be required. The large number of component parts presents challenges in the areas of integrated system management, fault tolerance, tuning, and diagnosis

of performance issues. These technical challenges are analogous to those the HPC community faced 15 to 20 years ago when the first MPP systems were developed. Once again, we have an opportunity to define requirements and develop solutions that effectively integrate our HPC file systems needs with COTS component technology. Delivered bandwidth, reliability, and cost are all critical issues that need to be addressed during the transition to petascale systems.

9) Reliability, availability, and serviceability (RAS) will need improvements in capability and functionality to support the ability to run millions of cores on a single large problem. We need the ability for integration/communication among the operating system, runtime system, application software, and parallel file system when failures occur. As noted above, the component part count for parallel file systems may drastically increase for petascale systems. The overall system will have to be highly resilient to failure of components.

10) Power and cooling requirements will place increasingly severe constraints on our ability to field petascale systems. While power and cooling are driving the move to multicore processors, we will be leveraging these commodity processors to create even larger integrated systems. Mainstream microprocessor vendors are designing with power awareness in mind, but future petascale and exascale advanced system initiatives must address overall system power limits in all aspects of hardware design and systems engineering.

APPENDIX 3.
Projected FYNSP Acquisition Plan

The following timeline shows the current ASC plan for platform deliveries over the life of the current FYNSP for Advanced, Capability and Capacity systems.

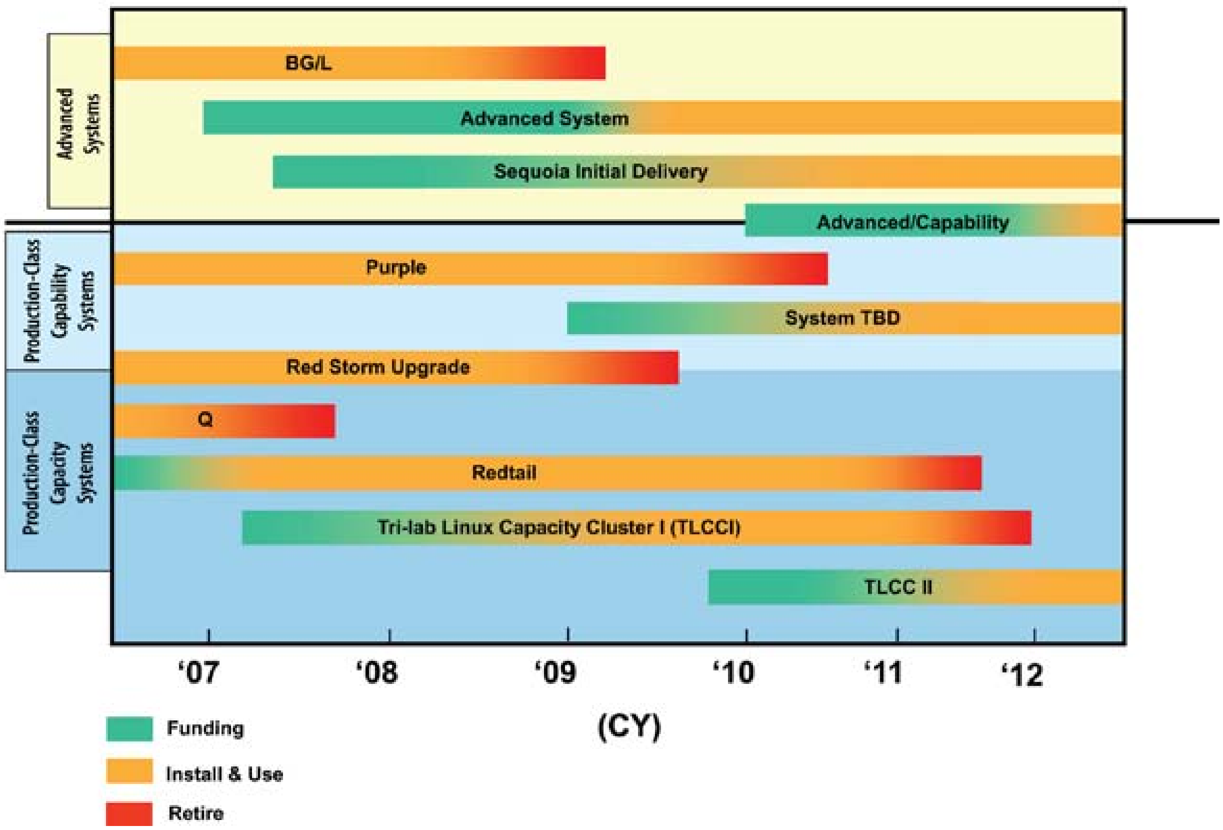

Figure 4. Timeline for platform deliveries

APPENDIX 4.
ASC Platform Governance Model

The following describes how, at a high level, ASC computing platforms will be allocated to meet the needs of the NNSA. Platforms are procured, sited, and operated by the NWC laboratories performing as executive agents of the ASC Program. These platforms are national assets focused on meeting the needs of the stockpile and supporting NNSA priorities. In keeping with these principles, ASC platform usage is governed by the following allocation procedures.

Capacity systems will normally be scheduled by the host laboratory to meet host laboratory priorities. However, the Program may direct the host to provide services to another laboratory to meet NNSA priorities.

Capability systems, upon reaching general availability, will be operated using a National User Facility model. While systems are in limited availability, friendly users will be solicited from across the tri-lab.

Advanced systems operations will be determined on a case-by-case basis, owing to the unpredictable nature of their associated operating environments. However, the program expects that some form of fair share scheduling would be employed that logically provides access to the tri-lab in equal proportions, not including system times required by the host for reliable and productive operations.

The Program will categorize platform procurements according to one of these three categories. As systems mature and requirements dictate, the Program may recategorize specific platforms, and may direct operations in a hybrid manner. Categories of systems (defined below) in the inventory at publication of this document are illustrative.

Platforms currently in the NWC inventory will be operated in the following manners:

Purple – Capability System

BG/L – Advanced System

Rhea/Minos/Lilac/UM/UV – Capacity Systems

Roadrunner Base – Capacity System

Lightning – Capacity System

Q – Capacity System

Red Storm – Hybrid Capability/Capacity System with 50% of system nodes used for tri-lab priorities for large-scale simulations

NWCC (Nuclear Weapons Computing Clusters) – Capacity System

TLCC (Tri-Laboratory Capacity Computing) – Capacity System

APPENDIX 5.
Glossary

Advanced Architectures
An ASC Program element that is focused on development of more effective architectures for high-end simulation and computing.

ASC
Advanced Simulation and Computing Program

ASC BG/L
An IBM system located at LLNL. In 2005, BlueGene/L was delivered as a 360 teraFLOPS system.

ASC Purple
An IBM system located at LLNL. In 2005, Purple was delivered as a 100 teraFLOPS system split between classified and open environments.

ASC Red Storm
A 40 teraFLOPS Cray system, located at SNL, delivered in FY2005, and upgraded to 124.4 teraFLOPS in 2006.

ASC Roadrunner Base
An IBM system located at LANL. In FY2006, a capacity computing system of 71 teraFLOPS was delivered to LANL for classified computing.

ASCI
Accelerated Strategic Computing Initiative

ASCI Blue Mountain
A Silicon Graphics, Inc. (SGI) system located at LANL. In 1998, ASCI Blue Mountain was installed as a 3.072 teraFLOPS computer system.

ASCI Blue Pacific
An IBM system located at LLNL. In 1998, ASCI Blue Pacific was installed as a 3.89 teraFLOPS computer system.

ASCI Q
A Compaq, now Hewlett-Packard (HP), system located at LANL. ASCI Q is a 20 teraFLOPS computer system, delivered in FY 2003.

ASCI Red
An Intel system located at SNL. ASC Red was the first general-purpose teraFLOPS platform in the world when it was installed in 1998 (1.872 teraFLOPS). Processor and memory upgrades in 1999 converted ASCI Red to a 3.15 teraFLOPS platform.

ASCI White
An IBM system located at LLNL. In 2000, ASCI White was installed as a 12.3 teraFLOPS supercomputer system.

BG/L
IBM BlueGene computer

capability/capacity systems
Terminology used to distinguish between systems that can run the most demanding single problems versus systems that manage aggregate throughput for many simultaneous smaller problems.

CCC
Capability Computing Campaign

CMOS
Complementary metal oxide semiconductor

COTS
Commercial-off-the-shelf

CPAC
Capability Planning Advisory Committee

CPU
Central processing unit

DOE
U.S. Department of Energy

DP
Defense Programs, one of the three major programmatic elements in NNSA.

DSW
Directed Stockpile Work, those SSP activities that directly support the day-to-day work associated with the refurbishment and certification of specific weapons in the nuclear stockpile.

exaFLOPS
Quintillion floating-point operations per second. ExaFLOPS is a measure of the performance of a computer.

FY
Fiscal Year. The U.S. Government's fiscal year runs from October 1 through September 30.

FYNSP
Future Years Nuclear Security Program

GF
GigaFLOPS

gigaFLOPS
Billion floating-point operations per second. GigaFLOPS is a measure of the performance of a computer.

HE
High energy explosive

HPC
High-Performance Computing

I/O
Input/output

KF
KiloFLOPS

LANL
Los Alamos National Laboratory, a prime contractor for NNSA, located in Los Alamos, New Mexico.

LEP
Life Extension Program

LLNL
Lawrence Livermore National Laboratory, a prime contractor for NNSA, located in Livermore, California.

LWK
Lightweight kernel

MF
MegaFLOPS

MTBF
Mean-time-between-failure

NNSA
National Nuclear Security Administration, a semi-autonomous agency within DOE

NWC
Nuclear Weapons Complex

petabyte
10^{15} bytes; 1,024 terabytes

petaFLOPS
Quadrillion floating-point operations per second. PetaFLOPS is a measure of the performance of a computer.

PF
PetaFLOPS

RAS
Reliability, Availability, and Serviceability

RRW
Reliable Replacement Warhead

science-based
The effort to increase understanding of the basic phenomena associated with nuclear weapons, to provide better predictive understanding of the safety and reliability of weapons, and to ensure a strong scientific and technical basis for future U.S. nuclear weapons policy objectives.

SFI
Significant Finding Investigation. An SFI results from the discovery of some apparent anomaly with the enduring stockpile. DSW Surveillance generally initiates an SFI.

SIMD
Simple-Instruction-Multiple-Data

SLEP
Stockpile Life Extension Program. SLEP is the DP element responsible for planning and execution of component and weapons refurbishments.

SNL
Sandia National Laboratories, a prime contractor for NNSA with locations in Albuquerque, New Mexico, and Livermore, California.

SSP
Stockpile Stewardship Program, DP's response to ensuring the safety, performance, and reliability of the U.S. nuclear stockpile.

SU
Scalable units

TCO
Total cost of ownership

teraFLOPS
Trillion floating-point operations per second. TeraFLOPS is a measure of the performance of a computer.

TF
TeraFLOPS

tri-lab
Refers to the three NNSA laboratories: LLNL, LANL, and SNL.

UGT
Underground testing

V&V
Verification and Validation. Verification is the process of confirming that a computer code correctly implements the algorithms that were intended. Validation is the process of confirming that the predictions of a code adequately represent measured physical phenomena.

ASC™

NNSA
National Nuclear Security Administration