



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Right Size Determining the Staff Necessary to Sustain Simulation and Computing Capabilities for Nuclear Security

D. J. Nikkel

September 24, 2010

Disclaimer


This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.



Right Size

Determining the Staff Necessary to Sustain Simulation and Computing Capabilities for Nuclear Security



NA-ASC-xxx-xx-xxxxxxxxxx

September 2010



INSIDE COVER PAGE

On the cover:

Text describing graphics used on cover.

Disclaimer

Auspices

SAND 2010-XXXXX

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

TITLE PAGE

Advanced Simulation and Computing

Right Size

Determining the Staff Necessary to Sustain Simulation and Computing Capabilities for Nuclear Security

Robert Meisner, Director

Daniel J. Nikkel, Technical Advisor

Advanced Simulation and Computing Program

NA-114

September 2010

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

Acknowledgments

This Right Size document is the product of many individual and team efforts. Our gratitude for their leadership and effort goes to all of them, and particularly to the following:

ASC Program Office at NNSA

Scott Doebling (LANL detailee)
Karen Pao (LANL detailee)
Brad Wallin (LLNL detailee)
Kim Yates (LLNL detailee)

ASC Executives

Michel McCoy (LLNL)
John Hopson (LANL)
James Peery (SNL)

Lawrence Livermore National Laboratory

Kimberly Budil
Brian Carnes
Kimberly Cupps
Desmond Pilkington
Gregory Pope
Teresa Quinn
James Rathkopf
Joseph Sefcik
Michael Zika

Los Alamos National Laboratory

Harold Armstrong
Mark Anderson
Bill Archer
Mark Chadwick
Robert Little
John Thorp
Cheryl Wampler
Robert Webster

Sandia National Laboratories

Sudip Dosanjh
Anthony Giunta
Joel Lash
John Noe
Constantine Pavlakos
Arthur Ratzel
David Rogers
David Womble

CONTENTS

Executive Summary	6
1 Introduction and Background	8
2 Defining the Right-Size Study	10
3 Findings	16
4 Validation.....	18
5 Conclusion	21
Appendix A. Technical Disciplines Identified in the Study	22
Appendix B. The Essential Role of Simulation in Stockpile Stewardship.....	24
Appendix C. Glossary.....	26

Executive Summary

The Advanced Simulation and Computing Campaign, herein referred to as the ASC Program, is a core element of the science-based Stockpile Stewardship Program (SSP), which enables assessment, certification, and maintenance of the safety, security, and reliability of the U.S. nuclear stockpile without the need to resume nuclear testing. The use of advanced parallel computing has transitioned from proof-of-principle to become a critical element for assessing and certifying the stockpile. As the initiative phase of the ASC Program came to an end in the mid-2000s, the National Nuclear Security Administration redirected resources to other urgent priorities, and resulting staff reductions in ASC occurred without the benefit of analysis of the impact on modern stockpile stewardship that is dependent on these new simulation capabilities. Consequently, in mid-2008 the ASC Program management commissioned a study to estimate the essential size and balance needed to sustain advanced simulation as a core component of stockpile stewardship. The ASC Program requires a minimum base staff size of 930 (which includes the number of staff necessary to maintain critical technical disciplines[·] as well as to execute required programmatic tasks) to sustain its essential ongoing role in stockpile stewardship.

To perform an estimate of the sustaining size of the program, hereinafter referred to as the “right-size” study, the study committee from the national weapons laboratories—Lawrence Livermore (LLNL), Los Alamos (LANL), and Sandia (SNL)—outlined three categories of information necessary to begin a comprehensive investigation:

- Identify the core disciplines required to cover the ASC mission.
- Categorize the years of training at a national weapons laboratory required to become an expert in an identified discipline (new recruit, approximately five years, or more than ten years).
- Identify the appropriate interaction between the weapons laboratories for each discipline (referred to as “interaction category,” or the 4Cs), in order to manage risk while eliminating unnecessary redundancy between laboratories.

Given the above input, the study committee completed a right-size estimate. They determined the required number of staff for each discipline, along with staff attributes and optimized inter-laboratory program balance. The study committee used rational, budget-independent methodologies to make its determination. The committee validated the estimate through various cross-checks, including comparison with historical numbers, the United Kingdom’s Atomic Weapons Establishment, and industry standard modeling estimates. Once the study committee completed and validated its estimate, the results and methodology of the study were assessed by six highly qualified domain experts external to ASC who individually concluded that the study’s methodologies were reasonable and

[·] The use of the term *discipline* in this document refers to a specialized technical area supporting the nuclear weapons program requiring skills gained through both academic training and laboratory experience (see Appendix A for examples).

the results were credible. In addition, the Carnegie Mellon Software Engineering Institute evaluated LLNL’s use of software development estimating models. They concluded that the models were applied appropriately, if somewhat conservatively, suggesting that laboratory estimates have some risk of being low.

The unanimous conclusion of the study and reviewers is that the current staff levels for the ASC Program are below what is needed to sustain science-based stockpile stewardship. The right-size study concluded that the ASC Program should be staffed at approximately 930 full-time staff members (compared to 782 employed at the beginning of FY11) across 85 disciplines. A summary of the study results is provided below. See Appendix A for a detailed breakdown of the study results.

Sector	Discipline Group	Staff Needed	Current Staff*
Simulation	Code generalist	15	10
	Hydrodynamics and materials	92	83
	Transport and plasmas	98	104
	Direct numerical simulation	32	13
	Engineering	105	84
	Simulation support	138	117
	Verification and validation	74	72
Simulation Totals		554	483
Computing	User support and outreach	63	38
	Center operations	194	160
	Production software development	29	30
	Management and business operations	16	7
	Facilities engineering	12	24
	Advanced architecture capabilities	62	40
Computing Totals		376	299
Overall Program Totals		930	782

* Current staff numbers refer to the beginning of FY11.

1 Introduction and Background

The Advanced Simulation and Computing (ASC) Program is a core component of the science-based Stockpile Stewardship Program (SSP), responsible for developing advanced simulation tools and computing hardware. The ASC Program mission, strategies, goals, and execution process are presented in a group of documents produced by the National Nuclear Security Administration (NNSA).^{1,2,3,4} The program, which began at about the same time as the moratorium on nuclear testing, has had two major phases: the Accelerated Strategic Computing Initiative (ASCI) and the ASC Program. ASCI focused on developing three-dimensional (3D) simulation tools and the high-performance machines and computing environments capable of running those tools and on establishing the proof-of-principle that advanced simulation could be used as part of a surrogate for new nuclear testing. The initiative's goal was to build an initial simulation capability for solving integrated multiphysics equations by developing computer platforms capable of 100 teraFLOPS performance, along with science and 3D integrated design codes capable of scaling to this level. The goals of the ASC Program are to support the continuing and time-constrained needs of stockpile stewardship and to enable development of a science-based predictive capability. In the ASC Program, Lawrence Livermore National Laboratory (LLNL) and Los Alamos National Laboratory (LANL) are responsible for the simulation tools associated with design and assessment of the nuclear explosive package. Sandia National Laboratories (SNL) is responsible for the tools associated with non-nuclear design and assessments. Competitive peer review is a cornerstone element of NNSA policy, the SSP, and the ASC Program. For more information on the ASC Program and its connection to the SSP, refer to Appendix B, The Essential Role of Simulation in Stockpile Stewardship.

Over the first decade of ASCI, staffing for the initiative increased to address the challenge of writing new simulation codes and developing state-of-the-art computer hardware. The fundamental goals of this initiative were met in 2005. The new computing capabilities and tools became essential elements of ongoing work in the SSP, and ASCI transitioned to the ASC Program. The use of advanced parallel computing transitioned from proof-of-principle to become a critical element for assessing and certifying

A 2008 benchmark study indicated LLNL staffing levels supporting code development were 65% of what is required based on industry standard estimation methods. A review of this study by the Carnegie Mellon Software Engineering Institute concurred and suggested that the estimated gap appeared to be conservative.

¹ *Advanced Simulation & Computing: The Next Ten Years*, Dimitri F. Kusnezov, 2004, NA-ASC-100R-04-Vol. 1-Rev. 0.

² *Advanced Simulation & Computing: Business Model*, Dimitri F. Kusnezov, 2004, NA-ASC-104R-05-Vol. 1-Rev. 0.

³ *Advanced Simulation & Computing Roadmap: National Nuclear Security Through Leadership in Weapons Science*, Dimitri F. Kusnezov and Njema Frazier, November 2006, NA-ASC-105R-06-Vol. 1-Rev. 0.

⁴ *A Platform Strategy for the Advanced Simulation and Computing Program*, Robert Meisner, August 2007, NA-ASC-113R-07-Vol. 1-Rev. 0.

the stockpile. As the initiative phase of the Nuclear Security Enterprise's simulation program came to an end, resources were redirected to other urgent priorities. The resulting ASC staff reductions occurred without the benefit of analysis of the impact on the new ways of assessing and certifying the stockpile enabled by these new simulation capabilities. A 2008 staffing benchmark study for software developers conducted at LLNL, using industry standard staff estimation methods, underscored concerns that program staff size had decreased too far. This study's conclusions, corroborated by a subsequent review by Carnegie Mellon's Software Engineering Institute (SEI), suggested software developer staffing at LLNL had dropped to a level creating increased risk for the program. In 2008, following multiple years of budget-driven staff reductions, NNSA sought to better understand its essential staffing requirements and undertook to estimate the size of the ASC Program necessary to sustain simulation and computing as a cornerstone of the science-based SSP. This estimate identified the size and balance of staff needed to sustain the current program over the long term, and this document describes the process for making this estimate and presents the results of the right-size study.

The ASC Program consists of five specific elements.² The *simulation* side of the program consists of three elements:

- Integrated Codes (IC).
- Physics and Engineering Models (PEM).
- Verification and Validation (V&V).

The *computing* side of the program consists of two elements:

- Computational Systems and Software Environment (CSSE).
- Facility Operations and User Support (FOUS).

The right-size study evaluated the simulation and computing sides of the program separately.

2 Defining the Right-Size Study

The scope of the ASC Program is given in the *ASC Strategy*¹ and *ASC Roadmap*³ documents, which present major objectives in both short-term elements necessary to meet continuing and time-constrained needs of stockpile stewardship, and longer-term elements necessary to ensure movement toward a science-based predictive capability that will enhance confidence in the simulation results. The right-size study was bounded by three inputs defined by the study committee, which the laboratories followed when preparing their estimates:

- Identify the core disciplines required to cover the ASC mission.
- Categorize the years of training at a national weapons laboratory required to become an expert in an identified discipline (new recruit, approximately five years, greater than ten years, or expertise in more than one discipline).
- Identify the appropriate interaction between the weapons laboratories for each discipline (distinguished among four “interaction categories,” or the 4Cs, described in Table 1) to ensure a balanced program and manage risks.

The second and third inputs above refine the estimation process beyond just identifying a single number for the various disciplines across the three laboratories. The second input acknowledges that to maintain a certain discipline with a necessary experience level, a “pipeline” of more novice staff is needed whose skills are being developed.[‡] The third input recognizes that not all technical disciplines must be staffed in equal depth at every laboratory. Total program staff requirements can be reduced by balancing efforts across the three laboratories using a method that manages risks.

Given these inputs, the laboratories estimated the number of staff for each discipline, then compared and contrasted them with other benchmarks. The estimation process is illustrated in Figure 1.

[‡] Many of the skills required to become an expert are not taught outside the laboratories, so new employees must undergo extensive on-the-job education and training by multiple mentors over extended time periods—5, 10, 15 years, or even longer. In addition to the number of experts needed, the laboratories must consider workforce demographics, which bear on both vacancies and the availability of mentors. Similar considerations apply to the pipeline of more novice staff, which also is subject to attrition.

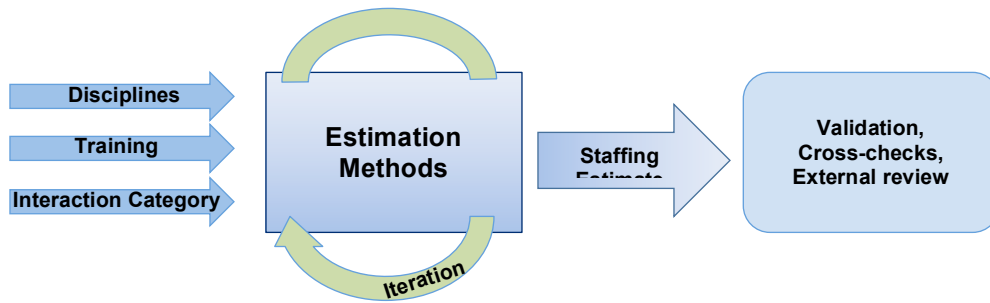


Figure 1. To determine the right size of the ASC Program, three inputs were defined, the staff necessary was estimated, and the results were cross-checked.

2.1 The Three Study Inputs

2.1.1 Listing the Detailed Disciplines

Core technical disciplines were identified to estimate manpower requirements by starting with a bottom-up approach, then subsequently placing disciplines in a hierarchy for ease of discussion.

Eighty-five disciplines were identified to cover the breadth of the ASC Program. For simulation, 69 disciplines were sorted into 7 discipline groups: code generalist, hydrodynamics and materials, transport and plasmas, direct numerical simulation, engineering, simulation support, and verification and validation. For computing, 16 disciplines were sorted into 6 discipline groups: user support and outreach, center operations, production software development, management and business operations, facilities engineering, and advanced architecture capabilities. See Appendix A, Technical Disciplines Identified in the Study, for a detailed listing of the discipline categories.

2.1.2 Identifying the Training Level for Each Discipline



For each discipline, the training level required within the program was identified by answering the question of how many years of training would be necessary for a productive staff member within each discipline to move from new recruit to expert. Some technical areas can capitalize on the skills and knowledge obtained right out of school, while others require years of mentoring by experts in program-specific activities to develop sufficient expertise. Thus, identifying the staff requirements for a given discipline was characterized in a more structured way than just a single number. The table in Appendix A indicates the training level for each discipline.

2.1.3 Determining the Interaction Category for Each Discipline

Within the ASC Program, managing risks within a limited budget (which includes ensuring an intellectually independent peer review capability) makes it essential to maintain staff to support independent efforts at different laboratories in some areas while also closely collaborating between laboratories in other areas to avoid unnecessary duplication of effort. The diversity of skills and degree of redundancy of skills required

between the laboratories is governed by the degree of difficulty and risk associated with particular technical goals. In particular, competitive efforts are required in areas with high technical risk that require strong peer review. The degree of laboratory interaction helps determine the staff necessary for each discipline. This is managed by placing activities into four interaction categories, collectively referred to as the 4Cs—competition, center of gravity, complementary, and commercial. Table 1 shows how the 4Cs were defined and provides sub-program examples.

Table 1. The four interaction categories (4Cs), their description, and examples within the ASC sub-programs.

	Interaction Category	Description	Simulation Examples	Computing Examples
High  Technical Risk  Low	Competition	Multiple significant efforts	Nuclear performance codes, plutonium equation of state	File systems, advanced architecture platforms
	Center of Gravity	One lab has a leading role, with key experts at other labs	Engineering codes, material damage, helium equation of state	Common computing environment, advanced systems research and development, archival storage
	Complementary	Moderate efforts at one or more labs integrating to achieve a common goal	Linear solver libraries, electron-ion coupling	Capacity computing platforms and environment
	Commercial	Off-the-shelf tools available	Configuration Management tools	Computer center operations, compilers

2.2 Estimating the Number of Staff for Each Methodology

The study committee, composed of a group of programmatic individuals from each laboratory, convened to develop methodologies for the right-size estimates within a program scope as defined by the *ASC Strategy* and *ASC Roadmap*. The committee was directed to guide its estimates with consideration of the following:

- The minimum size needed to execute the currently defined program.
- Budget independence and current mission (staff estimates driven by the mission requirements and not by the budget).
- Exclusion of significant future scope changes and new initiatives.

- The directed stockpile work mission requirements of independent design and assessment capabilities.
- Effective ASC collaborations with the Science and Engineering Campaigns to integrate the experimental data essential for validation and developing understanding within the Predictive Capability Framework (PCF).
- The need for a quantified predictive capability that requires significant increases in computation power (beyond exaFLOPS), thus demanding that ASC remain at the forefront of high performance computing (HPC) technology.
- The ability to successfully scale computer architectures and software to the required level.

It was noted that the number of required staff for the ASC Program does not scale with the number of weapons or weapon systems in the stockpile. Once the assumptions were agreed upon, the committee tasked the three laboratories to estimate their staff needs to deliver the requested mission product.

2.2.1 Right-Size Estimation Methods for Simulation

After collecting the input data, the study committee examined two independent methodologies for estimating the size of the detailed disciplines within the simulation components of the ASC Program (IC, PEM, and V&V). Method 1 used a training and development basis to establish a sustainable capability in each discipline. Method 1 recognizes the variations in levels of training of various experts required, the time it takes to develop them, and the necessity of maintaining a pipeline of staff in order to address attrition and sustain sufficient expertise at all required levels. Method 2 extrapolated current staff levels to a sustainable level to meet program goals. Method 2 starts with the current staff, the knowledge of where staffing is adequate and where there are shortages or excess, along with the known long-term program scope, and then adjusts current staffing in each discipline area to a level that could sustain ongoing work. The two methods give different perspectives on the right size of the program. If the estimate using Method 2 is greater than the estimate using Method 1, Method 2 estimated staffing is sufficient to maintain core disciplines. If the estimate using Method 1 is greater than the estimate using Method 2, it would imply core disciplines are understaffed at the Method 2 estimate level and will eventually be lost. The final right-size estimate applied both methods and iterated to a final result.

2.2.1.1 Method 1: Training and Development Basis

For Method 1, the study committee estimated the training requirements needed to sustain the disciplines within each subprogram. They started with the concept of a “code generalist”—an expert on most physics packages in IC and familiar with multiple codes. Maintaining several code generalists at each laboratory is essential for a healthy program, both for mentoring multidisciplinary staffs and for providing technical leadership to simulation development activities. Developing such a general expert requires a pool of more specialized experts, some of whom, over time, gain expertise in various physics

areas and codes. Building on multiple hierarchical levels of expertise and experience leads to an estimate of the number of staff ranging from entry level to experts at various levels of experience necessary to produce one code generalist. Because the program requires several of these highly experienced experts to carry out its mission, this methodology leads to an estimate of the total number of people needed by the program. This is referred to as the training and development based estimate. The pyramid representation in Figure 2 illustrates the training levels necessary, how it takes time to develop staff to the expert levels, and how there are fewer staff members with the necessary skills and experience at the higher levels in the pyramid. This is analogous to the process of developing a general officer in the military with broad training and experience from a larger pool of lower ranked officers who undergo a variety of leadership and training opportunities as they progress through the ranks.

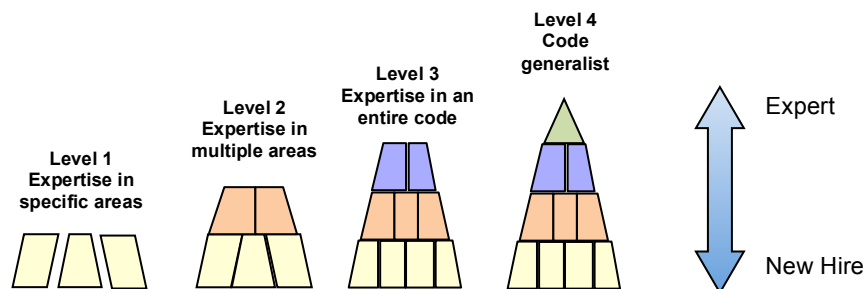


Figure 2. A pyramid illustrates the training levels for a given discipline. It takes significant time to develop necessary experts in critical areas from new hires.

Method 1 distinguished between a staff member as a new recruit, a staff member with approximately 5 years of experience, an expert with over 10 years of direct experience, and a seasoned expert in more than one discipline. The study committee estimated the maximum training levels required for each discipline area, the number of highest trained staff needed for each discipline, and the number of lower trained staff needed to maintain the required number at the highest trained level. Within the simulation side of ASC, six disciplines were identified as requiring top experts, while the majority of the disciplines required a mid-range level of expertise.

This method captured the hierarchical levels of staff required by the program (represented by the pyramid in Figure 2) and incorporated the ideas that:

- The availability of staff in levels three and four is essential to ensure innovation and ASC Program delivery.
- Staff move across program areas during a career.

2.2.1.2 Method 2: Sustainable Program Basis

For Method 2, the study committee identified both the scope of work and the core technical skills required by the mission in order to determine the staff levels necessary to maintain those disciplines and adhere to programmatic timelines for progress. Estimates considered the current staffing and extrapolated to the level at which the work could be

sustained in order to meet the defined program scope. This bottom-up approach began by identifying the categories of expertise required and the staff shortfall in each area, based on past management experience. The disciplines were divided into three categories according to the amount of experience required: long (more than 10 years), medium (approximately 5 years), and short (a new recruit up to 5 years). Eighty-five detailed disciplines were defined to cover the breadth of the ASC Program (see Appendix A). The Executive Summary shows these categories in higher-level discipline groups.

2.2.2 Right-Size Estimation Methods for Computing

To estimate the number of staff for the detailed disciplines needed within the computing component of the ASC Program (CSSE and FOUS), the study committee used a single method that was somewhat similar to Method 2 used in the simulation component, as described in the previous section.

The committee first identified the current staff in each discipline. Then they identified any changes required by the programmatic drivers (i.e., those identified in the *ASC Roadmap*). This step allowed for uncertainties due to new computing platforms and unknown future needs, with an internal review to remove double counting and implement the 4C methodology. The staff required for a discipline is a function of workload in that area, the degree of inter-laboratory interaction, and the required training. Finally, the laboratories cross-checked results and confirmed consistency of the 4Cs between laboratories, then compared numbers to industry standards for staffing HPC centers.

3 Findings

After completing the right-size estimates, results from both estimation methods were compared. The three laboratories demonstrated consistency where expected and determined that ASC staff levels should be approximately 930 full-time members. These numbers were broken down in detail across the technical disciplines and across the three laboratories. The Executive Summary shows a high-level breakdown of these numbers condensed into a few summary categories. The detailed numbers and discipline groups are provided in Appendix A.

Achieving the right-size staffing profile would require increasing some discipline groups by as much as 145% and decreasing some by 50%, with a net increase in program staff of approximately 19%.

For the simulation side of the ASC Program, the two estimation methods led to an overall estimate of staff requirements of approximately 554. As mentioned previously, Methods 1 and 2 represent distinctly different ways of looking at staffing the program, but the fact that these different approaches led to similar estimates provides evidence of the reasonableness of this final number.

The estimates for the computing side of the ASC Program, using a more traditional approach based on the size and complexity of the systems and the numbers of users being supported, gave an overall estimate of staff requirements of approximately 376. This estimate also included resources for software environments.

As described in Section 2.1.3, the study also identified appropriate levels of inter-laboratory interaction for each discipline as a way to balance the program, thereby eliminating unnecessary redundancy while minimizing risk. The staff numbers were distributed across the four interaction categories (see Table 1) associated with the various disciplines. As seen in Figure 3, the computing side of the program consists largely of complementary efforts, whereas the simulation side of the program, with more areas of higher risk, has more competitive and center of gravity staffing across the laboratories.




Figure 3. The right-size staff estimate distributed among the four interaction categories in the simulation and computing sides of the program shows how staff planning mitigates risk and minimizes redundancy.

The right-size estimate is only modestly greater than the staff in the pre-ASCI era of underground testing. For example, LANL's simulation team prior to ASCI was approximately 150 staff members, compared to the 186 estimated to be necessary by the right-size study. Stockpile stewardship's reliance on advanced simulation is much greater than in the pre-ASCI era, and the codes are required to be more dependable than ever before. This leads to increased staffing in the areas of V&V. Moreover, codes that run on the massively parallel modern computers are more complex than the serial and vectorized codes of the pre-ASCI area, thus increasing demands on code development, maintenance, and porting.

Although the estimates seem reasonable and track well against each other and the past programs, some reservations regarding the estimates remain. For example, the goal of developing "predictive simulation" on as-yet-undefined computing platforms is a difficult applied science problem that inherently has uncertainties in the path forward. Challenges may arise in the future, rendering the current staff estimates in certain technical areas inaccurate.

· In FY02, the ASC Program evolved from ASCI.

4 Validation

Several approaches were used to validate the methodologies used and conclusions reached in the right-size study. Before involving outside resources to review the estimation results, cross-checks were performed internally across the three laboratories. For example, using estimation Method 2 (Sustainable Program Basis) to cross-check estimation Method 1 (Training and Development Basis) results changed simulation estimates by less than eight percent and computing estimates by less than seven percent. After the cross-checks were complete, consistent numbers between the laboratories were achieved, with the physics laboratories (LANL, LLNL) simulation estimates agreeing within two percent and computing estimates for the three laboratories agreeing within five percent.

Results of the study were validated using industry standard methods, comparison with similar organizations, and independent external reviews by prominent experts.

In addition to the internal cross-checks, these additional steps were taken to validate the study:

- External domain expert reviews of the entire estimation process.
- Comparison review between the IC estimates and the Atomic Weapons Establishment (AWE) in the United Kingdom.
- Study of LLNL software developers required to support ASC codes using industry standard methodologies and subsequent review by Carnegie Mellon University's SEI.
- Computing resource estimates using comparable HPC centers outside NNSA.

All of these validation steps, which are described in more detail below, support the overall conclusion that the numbers estimated in the study are credible, and that current staffing is below the right size for the ASC Program.

4.1 Domain Expert Reviews

Six experts in HPC, simulation, and the NNSA mission—from DOE's Office of Science, the University of Michigan, the California Institute of Technology, Lawrence Berkeley National Laboratory, and the NNSA—were asked to review the resource estimation process and conclusions. The domain experts offering independent assessments were:

- Dr. Everet H. Beckner
Former NNSA Deputy Administrator for Defense Programs
- Mr. Stephen Guidice
Former NNSA Defense Programs Manager
- Ms. Barbara Helland
Advanced Scientific Computing Research Program
DOE Office of Science

- Prof. William Martin
Department of Nuclear Engineering and Radiological Sciences
University of Michigan
- Prof. Daniel I. Meiron
Fletcher Jones Professor of Applied & Computational Mathematics and
Aeronautics
California Institute of Technology
- Dr. Katherine Yelick
National Energy Research Scientific Computing Center Director
Lawrence Berkeley National Laboratory

The expert reviewers were told that the ASC Program performed a bottom-up study to estimate the number of staff required to perform its mission and were asked to give their opinion on that process and conclusions in light of the following questions:

- Were the methods used to determine the staff estimates reasonable?
- Were the number and balance of staff identified to achieve the assumed ASC mission reasonable?
- What major staffing issues may impact ASC in achieving its mission?

Each of the expert reviewers provided insights into the strengths and weaknesses of the estimate, but overall their comments concluded that the study’s methodologies were reasonable and that the findings were credible.

4.2 Comparison Review between ASC Integrated Codes and the Atomic Weapons Establishment

The NNSA and AWE approaches to managing their respective stockpiles are different, resulting in differently structured programs and staffing profiles. In an area where a direct comparison can be made—that of software development efforts on ICs—results showed that AWE staff totals for each major integrated design code project are similar to the ASC Program. This supports the conclusion that the right-size estimates of staff required for the code development portion of the study are reasonable.

4.3 Code Developer Benchmark and Subsequent Carnegie Mellon Review

One of the inspirations for the right-size study was a benchmark study of code developers at LLNL. This study was based on the complexity and lines of code in LLNL ASC codes using industry standard methods for estimating resources. It concluded that LLNL only has 65 percent of the code developers needed to successfully accomplish their mission. Subsequent to the right-size study, the LLNL code development benchmark study data were given to Carnegie Mellon University’s SEI for validation. The SEI is a federally funded research and development center that is well known for its expertise in computer science and software engineering project management.

SEI independently assessed the study and reported “the estimation techniques and tools used to develop the benchmark were appropriate and well documented.” They also

concluded that "... the benchmark accurately reflects a shortfall in staff supporting the ASC Program's efforts," and that "... the risks the program faces in supporting its mission is increasing significantly also due to staff age, not having the right type of personnel to support porting to the new computers the labs will be using, the complexity and age of the code base and the need for re-architecting the code base, and finally the time it takes to bring in and train new personnel."

4.4 Comparison with Computing Center Staffing outside of NNSA

ASC Program estimates for computing center personnel were compared to staff sizes at five other non-ASC centers from DOE and the Department of Defense. Four categories of data were analyzed: user support and outreach, center operations, software development, and management and business operations. The data collected were normalized rationally using parameters easily obtained (for example, the number of major computations systems, the FLOPS computing power, and the number of users at each facility). Data on the ASC computer centers showed operational efficiency consistent with, and at least as efficient as, that of external computer centers. Differences between the ASC centers and standard industry HPC centers were noted, including the complexity of running both unclassified and classified centers simultaneously with both production and advanced computers, support for remote classified computing, the scale of the total operation and leading-edge architecture, and self-maintenance that increases staff but reduces overall operating costs.

5 Conclusion

A quantitative determination of the appropriate size and balance of staff needed to perform the ASC mission began with a study at LLNL and was extended by ASC headquarters management to assess the entire program across all three national weapons laboratories. At each stage, recognized external experts have reviewed methodologies and analyses and judged them to be sound. The cross-checks and validation support the credibility of the conclusions of the right-size staffing estimate of approximately 930 for a sustainable program. The ASC staffing at the beginning of FY11, approximately 782 people, is roughly 19 percent below the right-size estimate. The ASC Program is currently meeting its mission objectives, but while the insidious effects of not having sufficient staff with appropriate critical skills do not manifest themselves in our current ability to accomplish the mission, it negatively impacts the future in complex ways that are difficult to accurately predict. If staff levels remain low, capabilities will eventually be lost as the staffing pipeline empties and attrition reduces the most experienced staff; progress will be slowed, and the ASC Program will be at increased risk of failing to meet some of its out-year commitments in support of the science-based SSP.

This right-size study provides a rational framework for rebuilding the ASC Program staff to a level that can sustain the critical role played by advanced simulation in stockpile stewardship and in the broader national security arena.

Appendix A. Technical Disciplines Identified in the Study

Detailed technical discipline areas identified in the right-size study, along with the estimated staff requirements, are shown below. The raw staff number estimates were supplied to within 0.5 staff members. (These numbers are rounded off and summarized in the Executive Summary.) The discipline areas are color-coded by the interaction category between the labs (4Cs), and the staff requirements numbers are color-coded by the training levels required. The particular identification of disciplines was influenced by the way the laboratories execute their mission and are not meant to indicate that distinct academic training is required for entry level staff in every discipline listed.

Color code for interaction categories

Competitive	Center of Gravity	Complementary	Commercial	Mixed
More Risk	←			Less Risk

Color code for required programmatic training level

Few years or new grad	~ 5 years	~ 10 or more years	Mixture of all
-----------------------	-----------	--------------------	----------------

	Discipline (color ~ interaction category)	Number (color ~ training level)
Simulation	Code Generalist	
	Code generalist	15
	Hydrodynamics and Materials	
	ALE hydro	24
	Eulerian hydro	8
	Equation of State	17.5
	Strength models	14.5
	Damage models	8.5
	High explosives	19.5
	Transport and Plasmas	
	Rad transport, Sn	10.5
	Rad transport, IMC	12
	Charged particle transport	5.5
	Neutronics transport, Sn	10
	Neutronics transport, MC	10.5
	ICF physics	5
	Opacities	10.5
	Plasmas	11.5
	Nuclear burn	7.5
	Nuclear data	15
	Direct Numerical Simulation	

Turbulence/mix	16.5
Multiscale physics	15
Engineering	
Thermal/fluid code development	10
Thermal/fluid model validation	2
Fire physics	3.5
Manufacturing processes	3
Reactive processes	3
Aerothermal/reentry	4.5
Non-continuum transport	4
Plasma physics	4
Solid mechanics code development	13
Solid mechanics model validation	2
Contact mechanics	2
Fracture and failure	6.5
Material model algorithms and implementation	2.5
Energy transfer/dissipation	3.5
Constitutive models	3
Microscale mechanics	1.5
Structural response	3

Fatigue and aging	1.5
Rad/electrical code development	8
Rad/electrical model validation	1
Lightning/electrostatic discharge	2
Neutron/ γ /x-ray material response	3.5
SGEMP/EMR/EMP	4.5
Component/device/system performance	7
Neutron/ γ /x-ray transport and deposition	0.5
Electrochemical aging	2.5
Electronic materials aging	3.5
Simulation Support	
Optimization/visualization/analysis	22.5
Code infrastructure	39
Numeric libraries	17
Software architecture and design	4
Setup – ALE	10
Setups – Eulerian	
Setups – Engineering	4
CM and automation tools	9.5
Management	19
Academic alliances	1
Administrative support	10
Verification and Validation	
Verification methods development	10
Validation methods development	2
UQ methods development	17
Diagnostic interpretation	3

Primary validation suite	7
Secondary validation suite	9
Safety validation suite	4
Diagnostic validation suite	2
Engineering validation suite	8
Small scale validation	3
UGT data reanalysis	5
Database validation	4
User Support and Outreach	
User support	41
Application environment, tools, support	22
Center Operations	
Systems integration and operation	42
File systems	28.5
Archival storage	24
Visualization and data analysis	21
Advanced display environments	6
Networking	23.5
Center operations	36
Security systems	12.5
Production Software Development	
System software	28.5
Management and Business Operations	
Management and business operations	16
Facilities Engineering	
Facilities engineering	12.5
Advanced Architecture Capabilities	
Performance modeling and system simulation	23
Petascale enablement	19.5
Advanced technologies and architecture	19

Computing

Appendix B. The Essential Role of Simulation in Stockpile Stewardship

The science-based SSP is an integrated program involving design analysis, stockpile surveillance, integrated and focused physics experiments, development of improved theoretical understanding of weapons physics and engineering, and simulation science. In the absence of nuclear testing, the foundation for maintaining the U.S. nuclear weapons stockpile is computational simulation that is informed and validated by experimental data from historical tests and the Science and Engineering Campaigns. The ASC Program develops for and provides to the SSP the computational hardware, software environments, application codes, theoretical models, and validation processes to underpin the use of simulations with confidence in assessing the current stockpile as well as future stockpile options.

ASC has, since its inception, been driven by the need to ensure the safety, reliability, and performance of the U.S. nuclear weapons stockpile without new nuclear testing. ASC's programmatic success results from it being a balanced program, driven by SSP mission deliverables, that provides all the elements—not just hardware—necessary for a simulation capability. This mission drive and balance has also differentiated it from many other HPC efforts.

The first decade of the program (the initiative phase) focused on developing computer platforms capable of 100 teraFLOPS performance along with science and 3D integrated design codes capable of scaling to this level. As these early capabilities were demonstrated, advanced simulation became a core part of stockpile stewardship and the annual assessment process. ASC is now chartered with the mandate to develop a predictive capability in the integrated design codes. With predictive codes (used with appropriate uncertainty quantification methodologies), simulations of systems diverging from the historical test database can be made with confidence, as weapons simulations will no longer be limited by being calibrated to tests in the archival database. The historical nuclear test data will still have great value, as it will be used for *validation* of capabilities, not *calibration* of simulations.

As computational simulation has become an essential component of stockpile stewardship, the capabilities enabled by ASC-class computing have put computational science on an equal footing with theoretical and experimental science as a tool for studying basic issues of weapons science and for scientific discovery. ASC platforms and tools are also the primary computational resources for ongoing, time-constrained work on the stockpile. The two major foci of the ASC Program strategy are:

- Meet the continuing and time-constrained needs of stockpile stewardship.

- Ensure progress toward the long-term goal of reduced dependence on phenomenology to enhance confidence—in other words, to move from calibrated simulations to predictive simulations.

To meet the continuing needs of stockpile stewardship, simulations using ASC codes running on ASC computers are heavily used by directed stockpile work in supporting life-extension programs and significant finding investigations, and for the annual assessment of systems in the stockpile. In addition to the integrated design codes used directly on stockpile system simulations, ASC also develops specialized physics and material property codes to study basic issues of weapons science and develop models and data used by the integrated design codes.

As the stockpile continues to age, the U.S. needs the ability to simulate the evolution of the state of the weapons and predict their performance in their current state. The second focus of the ASC Program, therefore, addresses the need to develop a predictive simulation capability with improved physical models as stockpile systems age further away from the historical test data. In the absence of nuclear testing, confidence in the simulations through quantifiable measures of the uncertainty in the predictions becomes essential. As was stated in the *ASC Roadmap*:

“... sustaining the testing moratorium requires that we transition to a point of sustainability at which our confidence in science-based simulations exceeds our confidence in simulations calibrated by underground test data.”

Whatever the future of the stockpile, it will be essentially different from what it is today (older, smaller, less diverse in systems but more diverse in individual differentiation), and future decisions regarding the stockpile will rely critically on ASC simulation capabilities in assuring the certainty to friend and foe of the readiness of the U.S. deterrent.

While at first glance one might think that sustaining the current program does not require the same kind of push as the first ASCI-scale threshold, the program still has major technical challenges and program deliverables, including advancing predictive capability needed by the SSP and developing the required exascale computing necessary to do that. Moreover, the challenges facing science-based stockpile stewardship are becoming steadily more difficult over time as the stockpile ages and the program is increasingly dependent on developing a validated predictive simulation capability.

Appendix C. Glossary

Advanced Simulation and Computing Program (ASC)

A cornerstone of the SSP that provides simulation capabilities and computational resources to support the ongoing work of stockpile stewardship.

Atomic Weapons Establishment (AWE)

The organization responsible for the United Kingdom's nuclear deterrent.

Computational Systems and Software Environments (CSSE)

The element of the ASC Program responsible for delivering and deploying the ASC computational systems and user environments via technology development and integration.

Facility Operations and User Support (FOUS)

The element of the ASC Program responsible for computer facility and operational support for reliable production computing and storage environments.

FLOPS

Acronym for Floating Point Operations per Second, a measure of computing performance. TeraFLOPS are a trillion (10^{12}) FLOPS, and exaFLOPS are a quintillion (10^{18}) FLOPS.

High Performance Computing (HPC)

HPC refers to the use of massively parallel supercomputers or computer clusters to solve advanced computational problems. Computers capable of teraFLOPS are counted as HPC systems.

Integrated Codes (IC)

The element of the ASC Program responsible for developing multiphysics simulation programs (codes) for use on massively parallel ASC platforms.

Integrated Design Codes

Large-scale computer programs for simulating design problems involving multiple physical phenomena.

Interaction Categories (4Cs)

Four categories describing a range of interaction between the laboratories in each discipline area, reflecting the need to manage risk by having some duplication of efforts while at the same time eliminating unnecessary redundancy.

Physics and Engineering Models (PEM)

The element of the ASC Program responsible for developing models of physics and material phenomena that get incorporated into the integrated design codes.

Predictive Capability

The ability to compute, with quantified confidence, the response of a system to external loading events, without prior knowledge of the outcome of the event. In the context of the nuclear weapons program, this means the ability to confidently simulate a weapon's performance, safety, and reliability with sufficient accuracy and with the uncertainties in the simulations identified and quantified, without the need to do new nuclear testing to calibrate simulations.

Predictive Capability Framework (PCF)

An NNSA Defense Programs integrated roadmap that lays out major technical goals among a variety of program elements, enabling integration, coordination, and prioritization of activities leading towards improved predictive capability.

Science and Engineering Campaigns

Sub-programs within NNSA Defense Programs responsible for the development of scientific and engineering tools necessary for improved predictive capability and assessment capabilities needed to support the stockpile in the absence of nuclear testing.

Software Engineering Institute (SEI)

A federally funded research and development center, operated by Carnegie Mellon University, with widely acknowledged expertise in software engineering.

Stockpile Stewardship Program (SSP)

A single, highly integrated technical program for maintaining the safety, security, survivability and reliability of the U.S. nuclear stockpile.

Verification and Validation (V&V)

The element of the ASC Program responsible for developing methods and metrics and performing assessments to ensure the accuracy of ASC simulations.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.