

# **FY07-08**

# **IMPLEMENTATION PLAN**

# **VOLUME 2, Rev. 0.5**

*October 18, 2006*



This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express 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 the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.





## Advanced Simulation and Computing

# FY07–08 IMPLEMENTATION PLAN Volume 2, Rev. 0.5

October 18, 2006

Approved by:

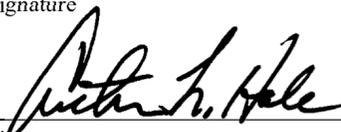
**Dimitri Kusnezov,**  
NNSA ASC Program Director

  
\_\_\_\_\_  
Signature

10/17/06

Date

**Arthur Hale,**  
SNL ASC Executive

  
\_\_\_\_\_  
Signature

9/13/06

Date

**Michel McCoy,**  
LLNL ASC Executive

  
\_\_\_\_\_  
Signature

9/21/07

Date

**John Hopson,**  
LANL ASC Executive

  
\_\_\_\_\_  
Signature

9/18/06

Date

**ASC Focal Point**  
**Dimitri Kusnezov**  
NA 114  
Tele.: 202-586-1800  
FAX: 202-586-0405  
dimitri.kusnezov@nnsa.doe.gov

**IP Focal Point**  
**Njema Frazier**  
NA 114  
Tele.: 202-586-5789  
FAX: 202-586-7754  
njema.frazier@nnsa.doe.gov

## Implementation Plan Contents at a Glance

Section No./Title	Vol. 1	Vol. 2	Vol. 3
I. Executive Summary			
II. Introduction			
III. Accomplishments			
IV. Product Descriptions			
V. ASC Level 1 and 2 Milestones			
VI. ASC Integration			
VII. ASC Risk Management			
VIII. Performance Measures			
IX. Budget			
Appendix A. Glossary			
Appendix B. Points of Contact			
Appendix C. 1.5.1.4.12 Alliance Centers			

# Contents

<b>I. EXECUTIVE SUMMARY.....</b>	<b>1</b>
<b>II. INTRODUCTION .....</b>	<b>2</b>
ASC Contributions to the Stockpile Stewardship Program .....	3
The National Work Breakdown Structure .....	6
<b>III. ACCOMPLISHMENTS FOR FY05–FY06 .....</b>	<b>9</b>
Computational Systems and Software Environment .....	9
Facility Operations and User Support .....	14
<b>IV. PRODUCT DESCRIPTIONS BY THE NATIONAL WORK BREAKDOWN STRUCTURE .....</b>	<b>17</b>
<b><i>WBS 1.5.4: Computational Systems and Software Environment .....</i></b>	<b><i>17</i></b>
WBS 1.5.4.1: Capability Systems .....	17
1.5.4.1.1 Purple (LLNL) .....	18
1.5.4.1.2 Systems Requirements and Planning (LANL) .....	18
1.5.4.1.4 Red Storm Capability Computing Platform (SNL) .....	20
1.5.4.1.5 Next Generation Capability Computing Platform (SNL) .....	22
WBS 1.5.4.2: Capacity Systems .....	22
1.5.4.2.2 Capacity Integration (LANL).....	23
1.5.4.2.3 Deployment of ASC Capacity Systems (SNL).....	23
WBS 1.5.4.3: Advanced Systems.....	24
1.5.4.3.1 BlueGene (LLNL).....	24
1.5.4.3.2 Advanced Systems Procurement and Supporting Technology Research and Development (SNL) .....	25
1.5.4.3.3 Blue Gene / P and BlueGene / Q Research and Development (LLNL) .....	26
1.5.4.3.4 Roadrunner Phase 3 Procurement (LANL) .....	26
WBS 1.5.4.4: System Software and Tools.....	26
1.5.4.4.1 System Software Environment for Scalable Systems (LLNL).....	28
1.5.4.4.2 Code Development and Performance Environment for Scalable Systems (LLNL) .....	29
1.5.4.4.3 System Software and Tools Development (LANL).....	30
1.5.4.4.4 Seamless User Environment (LANL) .....	31
1.5.4.4.5 Productive User Environment (LANL) .....	31
1.5.4.4.6 Application Readiness Team (LANL).....	31
1.5.4.4.7 Software and Tools for Scalability and Reliability Performance (SNL).....	32
1.5.4.4.8 ASC Programming Tools Project (SNL) .....	33
1.5.4.4.9 Software for Heterogeneous Systems (LANL).....	34
1.5.4.4.10 Application and System Performance Analysis (LANL).....	36
WBS 1.5.4.5: Input/Output, Storage Systems, and Networking.....	36
1.5.4.5.1 Archival Storage (LLNL) .....	38
1.5.4.5.2 File Systems (LLNL).....	39
1.5.4.5.3 Networking and Testbeds (LLNL).....	40
1.5.4.5.4 Archival Storage Services (LANL).....	40
1.5.4.5.5 HPSS and PSI Development (LANL) .....	41

1.5.4.5.6 Scalable Input/Output and Storage Systems (SNL).....43

1.5.4.5.7 Advanced Networking (SNL).....44

1.5.4.5.9 File Systems and I/O (LANL).....45

WBS 1.5.4.6: Pre- and Post-Processing Environments.....46

1.5.4.6.1 Scientific Visualization (LLNL) .....47

1.5.4.6.2 Scientific Data Management (LLNL).....48

1.5.4.6.3 Systems and Services Integration (LLNL).....48

1.5.4.6.4 Capacity Visualization (LANL) .....49

1.5.4.6.5 Capability Visualization (LANL) .....50

1.5.4.6.6 Scientific Visualization (SNL) .....51

1.5.4.6.7 Scientific Data Management (SNL) .....52

1.5.4.6.8 Infrastructure Deployment and Support (SNL).....53

1.5.4.6.9 Analysis Management and Prototyping (SNL).....53

1.5.4.6.10 Model Generation Tools (SNL).....54

**WBS 1.5.5: Facility Operations and User Support.....55**

WBS 1.5.5.1: Facilities, Operations, and Communications .....55

1.5.5.1.1 System Administration and Operations (LLNL) .....56

1.5.5.1.2 Software and Hardware Maintenance, Licenses, and Contracts (LLNL).....57

1.5.5.1.3 Computing Environment Security and Infrastructure (LLNL) .....57

1.5.5.1.4 Facilities Infrastructure and Power (LLNL).....58

1.5.5.1.5 Classified and Unclassified Facility Networks (LLNL) .....58

1.5.5.1.6 Wide-Area Classified Networks (LLNL).....59

1.5.5.1.7 Requirements and Planning (LANL) .....59

1.5.5.1.8 Roadrunner Base System Integration and Operations (LANL) .....59

1.5.5.1.9 Ongoing Network Operations (LANL) .....61

1.5.5.1.10 (Network) Infrastructure Integration (LANL) .....62

1.5.5.1.11 Ongoing Systems Operations (LANL).....62

1.5.5.1.12 Systems and Network Maintenance (LANL) .....64

1.5.5.1.13 Systems and Network Equipment (LANL) .....65

1.5.5.1.14 Ongoing Facilities (LANL) .....66

1.5.5.1.15 Power, Security, and Facilities Support (SNL).....67

1.5.5.1.16 Network Operations, Deployment, and Distance Computing (SNL).....68

1.5.5.1.17 Sandia Production Systems (SNL) .....71

1.5.5.1.18 Applications in Support of Manufacturing Production and Connectivity (Y-12) .....72

WBS 1.5.5.2: User Support Services .....72

1.5.5.2.1 Hotlines and System Support (LLNL) .....74

1.5.5.2.2 Integrated Computing Network Consulting, Training and Documentation, and External Computing Support (LANL).....74

1.5.5.2.3 Ongoing User Services (LANL).....75

1.5.5.2.4 Sandia Scientific Computing User Support (SNL).....77

WBS 1.5.5.3: Collaborations.....78

1.5.5.3.1 Program Support (LLNL) .....79

1.5.5.3.2 One Program/ Three Labs (SNL).....79

1.5.5.3.3 Program Support (LANL).....80

**V. ASC LEVEL 1 AND 2 MILESTONES .....81**

**VI. ASC INTEGRATION..... 103**

**VII. ASC RISK MANAGEMENT ..... 121**

**VIII. PERFORMANCE MEASURES ..... 125**

**APPENDIX A. GLOSSARY..... 129**

**APPENDIX B. POINTS OF CONTACT..... 131**

**APPENDIX C. 1.5.1.4.12 ALLIANCE CENTERS PROJECT (LLNL, LANL, SNL) .... 133**



## List of Tables

Table II-1. Defense Program Campaigns .....	5
Table V-1. Quick Look: Level 1 Milestone Dependencies .....	81
Table V-2. Quick Look: Level 2 Milestone Dependencies for FY07 .....	82
Table V-3. Quick Look: <i>Preliminary</i> Level 2 Milestone Dependencies for FY08 .....	87
Table VI-1. All Subprograms as Related to Integrated Codes.....	103
Table VI-2. All Subprograms as Related to Physics and Engineering Models .....	108
Table VI-3. All Subprograms as Related to Verification and Validation .....	111
Table VI-4. All Subprograms as Related to Computational Systems and Software Environment .....	115
Table VI-5. All Subprograms as Related to Facility Operations and User Support .....	117
Table VI-6. All Subprograms as Related to External Programs .....	119
Table VII-1. ASC's Top Ten Risks.....	121
Table VIII-1. ASC Performance Measures .....	125



# I. Executive Summary

---

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the safety and reliability of the U.S. nuclear stockpile. The SSP uses past nuclear test data along with current and future nonnuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program will require the continued use of current facilities and programs along with new experimental facilities and computational enhancements to support these programs.

The Advanced Simulation and Computing Program (ASC)<sup>1</sup> is a cornerstone of the SSP, providing simulation capabilities and computational resources to support the annual stockpile assessment and certification, to study advanced nuclear-weapons design and manufacturing processes, to analyze accident scenarios and weapons aging, and to provide the tools to enable Stockpile Life Extension Programs (SLEPs) and the resolution of Significant Finding Investigations (SFIs). This requires a balanced resource, including technical staff, hardware, simulation software, and computer science solutions.

In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. In its second decade, ASC is focused on increasing its predictive capabilities in a three-dimensional simulation environment while maintaining the support to the SSP. The program continues to improve its unique tools for solving progressively more difficult stockpile problems (focused on sufficient resolution, dimensionality and scientific details); to quantify critical margins and uncertainties (QMU); and to resolve increasingly difficult analyses needed for the SSP. Moreover, ASC has restructured its business model from one that was very successful in delivering an initial capability to one that is integrated and focused on requirements-driven products that address long-standing technical questions related to enhanced predictive capability in the simulation tools.

ASC must continue to meet three objectives:

- **Objective 1. Robust Tools.** Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs, LEPs, annual assessments, and evolving future requirements.
- **Objective 2. Prediction through Simulation.** Deliver validated physics and engineering tools to enable simulations of nuclear-weapons performances in a variety of operational environments and physical regimes and to enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
- **Objective 3. Balanced Operational Infrastructure.** Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work (DSW) and SSP needs for capacity and high-end simulation capabilities.

---

<sup>1</sup> In FY02 the Advanced Simulation and Computing (ASC) Program evolved from the Accelerated Strategic Computing Initiative (ASCI).

## II. Introduction

---

The ASC Program supports the National Nuclear Security Administration's (NNSA's) long-term strategic goal of Nuclear Weapons Stewardship: *"ensure that our nuclear weapons continue to serve their essential deterrence role by maintaining and enhancing the safety, security, and reliability of the U.S. nuclear weapons stockpile."*<sup>2</sup>

In 1996, ASCI—the Accelerated Strategic Computing Initiative—was established as an essential element of the SSP to provide nuclear weapons simulation and modeling capabilities.

In 2000, the NNSA was established to carry out the national security responsibilities of the Department of Energy, including maintenance of a safe, secure, and reliable stockpile of nuclear weapons and associated materials capabilities and technologies.

Shortly thereafter, in 2002, ASCI matured from an initiative to a recognized program and was renamed the Advanced Simulation and Computing (ASC) Program.

Prior to the start of the nuclear testing moratorium in October 1992, the nuclear weapons stockpile was maintained through (1) underground nuclear testing and surveillance activities and (2) "modernization" (i.e., development of new weapons systems). A consequence of the nuclear test ban is that the safety, performance, and reliability of U.S. nuclear weapons must be ensured by other means for systems far beyond the lifetimes originally envisioned when the weapons were designed.

NNSA will carry out its responsibilities through the twenty-first century in accordance with the current Administration's vision and the Nuclear Posture Review (NPR) guidance. NNSA Administrator Ambassador Brooks summarized<sup>3</sup> the NNSA objectives for SSP as follows:

*"We will continue to lead the way to a safer world through the deep reductions in nuclear forces codified by the Moscow Treaty, through Nunn-Lugar and other cooperative threat reduction efforts, and through other actions. At the same time, although conventional forces will assume a larger share of the deterrent role, we will maintain an effective, reliable, and capable—though smaller—nuclear force as a hedge against a future that is uncertain and in a world in which substantial nuclear arsenals remain. Our ongoing efforts to reduce the current stockpile to the minimum consistent with national security requirements, to address options for transformation of this smaller stockpile, and to create a responsive nuclear weapons infrastructure are key elements of the Administration's national security strategy..."*

A truly responsive infrastructure will allow us to address and resolve any stockpile problems uncovered in our surveillance program; to adapt weapons (achieve a capability to modify or repackage existing warheads within 18 months of a decision to enter engineering development); to be able to design, develop, and initially produce a new warhead within three to four years of a decision to do so;<sup>4</sup> to restore production

---

<sup>2</sup> NNSA Strategic Plan, page 8.

<sup>3</sup> Speech presented to the Heritage Foundation Conference: *U.S. Strategic Command: Beyond the War on Terrorism*, May 12, 2004.

<sup>4</sup> While there are no plans to develop new weapons, gaining the capability is an important prerequisite to deep reductions in the nuclear stockpile.

capacity to produce new warheads in sufficient quantities to meet any defense needs that arise without disrupting ongoing refurbishments; to ensure that services such as warhead transportation, tritium support, and other ongoing support efforts are capable of being carried out on a time scale consistent with the Department of Defense's ability to deploy weapons; and to improve test readiness (an 18-month test readiness posture) in order to be able to diagnose a problem and design a test that could confirm the problem or certify the solution (without assuming any resumption of nuclear testing).

Additionally, the NPR guidance has directed that NNSA maintain a research and development and manufacturing base that ensures the long-term effectiveness of the nation's stockpile and begin a modest effort to examine concepts (for example, Advanced Concepts Initiatives, including the Robust Nuclear Earth Penetrator) that could be deployed to further enhance the deterrent capabilities of the stockpile in response to the national security challenges of the twenty-first century.

The ASC Program plays a vital role in the NNSA infrastructure and its ability to respond to the NPR guidance. The program focuses on development of modern simulation tools that can provide insights into stockpile problems, provide tools with which designers and analysts can certify nuclear weapons, and guide any necessary modifications in nuclear warheads and the underpinning manufacturing processes. Additionally, ASC is enhancing the predictive capability necessary to evaluate weapons effects, design experiments, and ensure test readiness.

ASC continues to improve its unique tools to solve progressively more difficult stockpile problems, with a focus on sufficient resolution, dimensionality, and scientific details, to quantify critical margins and uncertainties (QMU), to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW provides requirements for simulation, including planned SLEPs, stockpile support activities that may be ongoing or require short-term urgent response, and requirements for future capabilities to meet longer-term stockpile needs. Thus, ASC's advancing leading-edge technology in high-performance computing and predictive simulation meets these short- and long-term needs, including the annual assessments and certifications and SFIs. The following section lists past, present, and planned ASC contributions to meet these needs.

## **ASC Contributions to the Stockpile Stewardship Program**

**In FY96**, ASCI Red was delivered. Red, the world's first teraflops supercomputer, was upgraded to more than 3 teraflops in FY99 and was retired from service in September 2005.

**In FY98**, ASCI Blue Pacific and ASCI Blue Mountain were delivered. These platforms were the first 3-teraops systems in the world and have both since been decommissioned.

**In FY00**, ASCI successfully demonstrated the first-ever three dimensional (3D) simulation of a nuclear weapon primary explosion and the visualization capability to analyze the results; ASCI successfully demonstrated the first-ever 3D hostile-environment simulation; and ASCI accepted delivery of ASCI White, a 12.3-teraops supercomputer, which has since been retired from service.

**In FY01**, ASCI successfully demonstrated simulation of a 3D nuclear weapon secondary explosion; ASCI delivered a fully functional Problem Solving Environment for ASCI White; ASCI demonstrated high-bandwidth distance computing between the three national laboratories; and ASCI demonstrated the initial validation methodology for early primary behavior. Lastly, ASCI completed the 3D analysis for a stockpile-to-target sequence (STS) for normal environments.

**In FY02**, ASCI demonstrated 3D system simulation of a full-system (primary and secondary) thermonuclear weapon explosion, and ASCI completed the 3D analysis for an STS abnormal-environment crash-and-burn accident involving a nuclear weapon.

**In FY03**, ASCI delivered a nuclear safety simulation of a complex, abnormal, explosive initiation scenario; ASCI demonstrated the capability of computing electrical responses of a weapons system in a hostile (nuclear) environment; and ASCI delivered an operational 20-teraops platform on the ASCI Q machine.

**In FY04**, ASC provided simulation codes with focused model validation to support the annual certification of the stockpile and to assess manufacturing options. ASC supported the life-extension refurbishments of the W76 and W80, in addition to the W88 pit certification. In addition, ASC provided the simulation capabilities to design various nonnuclear experiments and diagnostics.

**In FY05**, ASC identified and documented SSP requirements to move beyond a 100-teraops computing platform to a petaFLOPS-class system; ASC delivered a metallurgical structural model for aging to support pit-lifetime estimations, including spiked-plutonium alloy. In addition, ASC provided the necessary simulation codes to support test readiness as part of NNSA's national priorities.

**By FY06**, ASC will deliver the capability to perform nuclear performance simulations and engineering simulations related to the W76/W80 LEPs to assess performance over relevant operational ranges, with assessments of uncertainty levels for selected sets of simulations. The deliverables of this milestone will be demonstrated through 2D and 3D physics and engineering simulations. The engineering simulations will analyze system behavior in abnormal thermal environments and mechanical response of systems to hostile blasts. Additionally, confidence measures and methods for uncertainty quantification will be developed to support weapons certification and QMU Level 1 milestones.

**By FY07**, ASC will support the completion of the W76-1 and W88 warhead certification, using quantified design margins and uncertainties; ASC will also provide a robust 100-teraOPS-platform production environment supporting DSW and Campaign simulation requirements. This will be augmented by the 360-teraFLOPS ASC BlueGene/L system, which provides additional capability for science campaigns.

**By FY08**, ASC will deliver the codes for experiment and diagnostic design to support the CD-4 approval on the National Ignition Facility (NIF). An advanced architecture platform capable of sustaining a 1-petaFLOPS benchmark will be sited at Los Alamos National Laboratory (LANL).

**By FY09**, a modern baseline of all enduring stockpile systems, using ASC codes, will be completed.

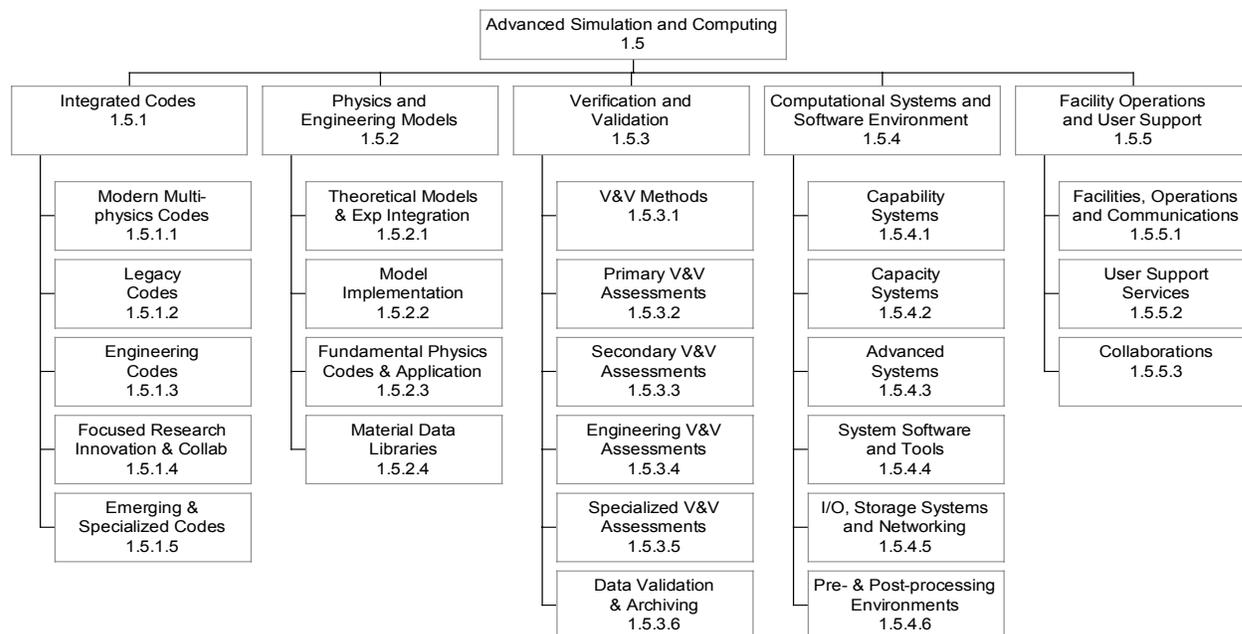
**In FY10 and beyond**, ASC will continue to deliver codes for experiment and diagnostic design to support the indirect-drive ignition experiments on the NIF and will continue to improve confidence and response time for predictive capabilities to answer questions of vital importance to the SSP.

**Table II-1. Defense Program Campaigns**

<b>Campaign Number</b>	<b>Campaign Title</b>
C1	Primary Assessment Technology and Test Readiness
C2	Dynamic Materials Properties
C3	Advanced Radiography
C4	Secondary Assessment Technology
C5	Enhanced Surety
C6	Weapon Systems Engineering Assessment Technology
C7	Nuclear Survivability
C8	Enhanced Surveillance
C9	Advanced Design & Production Technologies
C10	Inertial Confinement Fusion Ignition and High Yield Campaign
C11	Advanced Simulation and Computing
C12	Pit Manufacturing and Certification Campaign
C15	Nonnuclear Readiness
C16	Materials Readiness
C18	Engineering Campaigns Construction Activities
C19	Advanced Design & Production Technologies Readiness

## The National Work Breakdown Structure

ASC's program structure is based on the new national work breakdown structure (nWBS), described in the ASC Business Model (NA-ASC-104R-05-Vol.1-Rev.5).



### Sub-Programs

As the chart visualizes, ASC is divided into five sub-programs:

- Integrated Codes
- Physics and Engineering Models
- Verification and Validation
- Computational Systems and Software Environment
- Facility Operations and User Support

The first three sub-programs focus on improved models in the modern codes, delivery of validated tools, and response to SSP issues (for example, SFIs, LEPs, annual assessments). Key drivers are to improve the confidence in prediction through simulations; to calculate, measure, and understand the uncertainty in the predictions; and to ensure rapid delivery of simulation capabilities to the SSP.

The fourth sub-program, Computational Systems and Software Environment, ensures the development and deployment of a computing environment needed for all ASC-deployed platforms: capability, capacity, and advanced systems.<sup>5</sup> Not only is this sub-program responsible for related research and technology development, but it is also responsible for planning, procurement, and quality control activities.

<sup>5</sup> The ASC Program is in transition for current platforms. Future platforms will follow the Capital Acquisition Management process identified in the NA-10 *Program Management Manual*.

The fifth, and last, sub-program, Facility Operations and User Support, provides operational support for production computing and storage, user support services, and collaborative research opportunities with educational institutions, as well as programmatic support across the ASC program.

### **Product Deliverables**

The Product deliverables are described at level 4 and span the full-scope of the program in the context of the nWBS. They describe what the Laboratories expect to provide to a given Product as a result of their activities.

Deliverables can, but do not necessarily, contribute to level 2 milestones chosen in a given fiscal year. Deliverables that do directly contribute to level 2 milestones for the fiscal year will be reviewed in the context of level 2 milestone reviews.



### III. Accomplishments for FY05–FY06

---

ASC accomplishments from Quarter 4, fiscal year 2005, through quarter 3, fiscal year 2006, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

HQ is pleased to highlight the outstanding achievements of the Defense Programs Contractors.

#### Computational Systems and Software Environment

##### LLNL Accomplishments for Computational Systems and Software Environment

At LLNL, the major CSSE accomplishments related to formal L2 and significant L3 milestones are:

- The ASC Purple hardware deliveries were formally accepted and integrated into full systems. In order to supply ASC Level 1 Alliance support, 6.6 teraFLOPS (108 nodes) have been left on the unclassified side, while the remaining 93.4-teraFLOPS system is operational on the classified side and in limited availability use for stockpile requirements.
- The ASC BlueGene/L system was fully integrated and moved to the classified side for ASC production science use in February 2006. Applications running on the system during final system test, before moving to classified operations, included four finalists for the prestigious Gordon Bell Prize and established BlueGene/L as Number 1 on the Top 500 list with a performance of 280.6 teraFLOPS on the Linpack benchmark. The Gordon Bell Prize winning application, a 3D molecular-dynamics code, ddcMD, sustained 101.7 teraFLOPS for more than seven hours. Even faster performance runs of 207 teraFLOPS by a first-principles molecular-dynamics application, Qbox, were recently achieved and should be in position for this year's Gordon Bell Prize.
- LLNL successfully deployed a new security infrastructure within the classified computing environment. The key components of this new infrastructure are a Kerberos authentication service, a security registry based on the Lightweight Directory Access Protocol (LDAP), and an inter-site directory service, which are currently operating as production services. These services support login access and identity management of user accounts on Purple and BlueGene/L, for both local and inter-site users. Additionally, the new ASC security infrastructure is utilized by Sandia NWIE and WebFileShare web-based inter-site applications.
- The new LLNL multi-TF Gauss visualization cluster used to support BlueGene/L data analysis was successfully installed, integrated, and deployed for customer use, and recently moved in February 2006 to the classified side. A formal release of the VisIt tool incorporating hardware-accelerated parallel rendering was completed, as were formal releases of the Hopper and Telepath tools that incorporate new user-required data and resource management functionality. The Terascale Simulation Facility (TSF) classified Tilden Visualization Theater (formerly the Advanced Simulation Laboratory) deployment was also completed and has been used for several high-level reviews and meetings.

In other accomplishments at LLNL, a common resource manager (SLURM) is now deployed across all Linux systems and AIX platforms. SLURM was originally developed for Linux systems but was ported to Purple and all other LC AIX platforms in FY2006, replacing LoadLeveler. The high performance storage system (HPSS) development collaboration completed and released HPSS R6.2, which, among other things, completely removes archive dependencies on distributed computing environment (DCE) infrastructure. LLNL deployment teams successfully deployed this release into production in both classified and unclassified computer centers in support of production operation of ASC BlueGene/L and Purple platforms. The beta release of Open | SpeedShop, a PathForward product, has been installed on the Linux systems.

### **LANL Accomplishments for Computational Systems and Software Environment**

The Roadrunner Project was initiated, based upon a Congressional earmark in FY06 to bring a significant source of capacity computing resource to LANL in FY07 and to form the basis for an accelerated platform technology for applications and to set a 1-petaFLOPS Linpack benchmark. In FY06, the project was formed, a request for proposal was issued, and the contract was competed using standard procurement practices. At the time of writing, the contract is near award and announcement. It is anticipated that this will be completed and that hardware delivery for the base system will be completed in FY06.

The Crestone project codes were instrumented to report performance data. A canonical problem was selected and baseline performance measures were made on Lightning (commodity hardware Linux capacity cluster) and Q (custom-built capability cluster). Performance measures and comparisons were integrated into the build and test system so that performance can be tracked over time with changes to the codes. Measurements will direct developers and optimizers to places in the codes where performance improvements can be obtained.

The ViewMaster visualization cluster was delivered and deployed in the secure for production visualization use. There are direct office connections to ViewMaster for many designers and scientists. These users see significant rendering and application performance improvements over the existing SGI-based visualization infrastructure. Visualization applications have been ported and all local and wide-area network connections are up and running. Users are pleased with the quality of stereo and the vivid colors that can be seen using ViewMaster's digital fiber links instead of the old lossy analog links. The initial evaluation of the IBM CELL architecture as a hardware platform for visualization is complete. The CELL architecture and the GPU both show great promise for extremely fast visual analysis of complex datasets. We also extended our distance/streaming visualization software infrastructure to prioritize the processing and display of unstructured datasets.

The expert visualization staff continued providing quality support to users. In the weapons program, several important discussions arose from visualizations produced with the help of the expert visualization staff. Of special note for this year is the staff's assistance to the Reliable Replacement Warhead (RRW) and Verification and Validation (V&V) projects in using the PowerWall and CAVE, and in particular, the multi-lab working meetings hosted in the PowerWall theater, which resulted in RRW design issues being noticed for the first time and corrected. Other production visualization staff members are helping to stand up the new Viewmaster cluster and have helped decommission one of the old visualization systems safely and securely.

The Performance Analysis Lab team completed an in-depth analysis of Livermore's Purple system. They showed how the system configuration was limiting the performance to about half of its capability. Some of these issues were related to system

noise, similar to the problems discovered on Q and others are related to how the system is architected. They indicated that they expected close to a 100-percent improvement in application performance by correcting the system noise matters, but they could not change the architectural matters. The architectural matters are more significant for bandwidth-sensitive applications. The analysis was documented in a technical report.

PaScalBB's (Parallel Scalable Back Bone Concept) implementation was augmented this year to place disk blades in different subnets but logically form a single file system. File system clients route data through the proper subnet to reach the destination disk blades most efficiently. This approach allows for increased load balance, failover, and scalability. The file system can grow increasingly larger by adding more subnets, switches, input/output (I/O) nodes, and disk blades. PaScalBB support for Infiniband and 10 gigE was also demonstrated. File Transfer Agents were also added to the system this year to offload the work of moving files from the global parallel file system (GPFS) to the archive and allow the systems' compute nodes to be used for doing calculations rather than archiving files.

BTime software provides accurate time synchronization between Linux cluster nodes. It was released with a kernel module for Linux 2.6.x kernels, which extracts kernel timekeeping variables without changing system calls. This will allow an agile way to ensure that the time on Linux cluster nodes stays synchronized within few tens of microseconds of each other.

The OpenMPI team has begun working with the X division codes to ensure OpenMPI meets their functional and performance requirements. OpenMPI V1.1 was released for testing to application developers in June 2006. A noteworthy feature of V1.1 is the run-time selection of an optimized suite of collectives. A specification for a next-generation framework of automatically tuning collectives is being finalized, with implementation expected to begin soon.

OpenRTE successfully demonstrated a prototype capability for launching jobs on remote systems. It is currently working on a next-generation release that will support multi-cluster operations and is finalizing collaborative relationships to extend the system to support a number of new environments such as the NASA/NSF National Virtual Observatory's computational and control grid.

A small effort was begun to look at an option for a file system being able to scale to tens of thousands of clients and to be stateless, enabling process migration. PVFS2 (Parallel Virtual File System, V2) was selected. It was modified to provide an interface to Panasas' GPFS and show that files could be accessed via the PVFS2 or standard Linux POSIX interfaces. Scalability investigations are planned for the coming year.

Workflow models were developed based on case studies of Directed Stockpile Work (DSW) being performed by weapons designers using products from the ASC Shavano and Crestone projects. A baseline definition of CSSE project metrics was delivered that recommends specific measures for the program to address issues and impediments in the ASC production workflow. The workflow data contradict current perceptions held in the Program, namely that Setup dominates in the primary designer's workflow, and that Execution dominates in the secondary designer's workflow. Instead, six case studies suggest that Execution is consistently the more dominant factor in the elapsed time spent using ASC products. Setup dominates end-user time, however, leading to the conclusion that the ratio of Setup versus Execution time is a productivity issue rather than the absolute magnitude of Setup time itself. Also, the workflow baseline surfaces previously un-quantified data on the impact that non-predictive processes (for example, uncertainty quantification) and administrative tasks (for example, execution monitoring) have on DSW productivity. These effects have not been measured before yet they

constitute significant portions of mission-critical workflow. Future work may include automation of metric gathering, correlating the metric to overall productivity, assessing whether or not a reduction in the metric will have a significant impact on overall productivity and if the return on investment to reduce the metric is worthwhile, and showing changes of productivity over time and the sources of them.

BProc V4 was ported to the Linux 2.6.x 64-bit kernel. It was used on the 64-bit software stacks for the Lightning/Bolt and Flash/Gordon systems at LANL.

### **SNL Accomplishments for Computational Systems and Software Environment**

Red Storm successfully completed two level 2 milestones related to limited availability. The system achieved limited availability (MS #17) and the tri-lab problem-solving environment achieved limited availability (MS #18), as well (Q4FY05). Red Storm has made great strides since entering limited availability. It has shown impressive performance, scalability, and efficiency. This assertion has been validated by HPL and high-performance computing (HPC) benchmarks and by application usage. Applications have been running in the 5000-node regime since January. Attention is given to the scalability of the applications to ensure the processors are effectively utilized. Reliability continues to improve as well. The mean time between system-wide interrupt is in the desired range. Users from Sandia, Los Alamos, and Livermore, using both classified and unclassified portions of the machine, are routinely running production jobs in preparation for NW milestones.

The lightweight kernel operating system, Catamount, is the cornerstone of the Red Storm run time environment. During this past year, support was added for dual core Opteron processors. Like ASCI Red, a virtual node implementation was chosen so that no changes are required to the application. It can become immediately useful without any coding effort by application developers. Each processor core operates as a separate node, each utilizing half of the memory on the physical node and sharing the network interface. This feature can double the amount of concurrent work being done on the machine. Successful completion of this work was critical so that the system software would support the planned hardware upgrade in Q4FY06.

The Design-through-Analysis (DTA) Realization Team completed the initial release of the DTA environment on the Sandia Restricted Network (Q4FY05). The DTA environment comprises a federation of analysis setup and post-processing tools that interoperate through a common metadata standard. These tools help analysts rapidly complete tasks such as CAD model geometry editing, mesh building and editing, application of mesh parameters (for example, boundary conditions and materials information), input deck creation, job submission, post-processing and visualization of analysis results, and archiving of analysis artifacts.

The DTA environment was deployed on the Sandia Classified Network and the tools became available for application to classified weapon analysis problems (Q3FY06).

DTA tool capability improvements include the following: CUBIT released version 10.1 with new features focused on creating an overlay topology for solid models, a new algorithm for joining parts of a complex assembly to facilitate a conformal mesh as well as development of new tools for editing and improving an existing mesh; the analysis process coordinator implemented a project navigator to facilitate local and remote file management and a configuration management system to help analysts manage their many analysis artifacts; materials/WISDM has incorporated MIL-5 handbook and ASM material properties data as well as an AWE materials property database; and ASETS, the job submission tool, has implemented support for all targeted Linux cluster variants (Nuclear Weapons Compute Clusters (NWCC)/Institutional Computing

Cluster/ESHPC) as well as support for CALORE (thermal analysis), PRESTO (non-linear structural mechanics), and SALINAS (structural dynamics).

Sandia visualization software developers used ParaView to complete Milestone #1313, *Provide visualization for capability calculations on Red Storm*. Sandia-developed scalable capabilities within ParaView leverage the unique resources of the Red RoSE cluster to provide the only visualizations of many of the largest of the Red Storm simulations, including two of the initial science runs. In addition, Sandia delivered multiple advanced visualization capabilities within ParaView, such as parallel volume rendering and providing native support for P and H-adaptive meshes (Q4FY05).

ParaView now delivers its scalable performance and advanced visualization capabilities directly to the user's desktop through an advanced image compression and delivery system developed by Sandia. ParaView, running on the Red RoSE cluster, delivered a record speed of 8 billion triangles per second to the desktop. This record and other Sandia-developed capabilities have been highlighted in two recent NVidia/Sandia press releases.

The Red RoSE Cluster and DSA-HPSS achieved the Limited Availability Milestone. Red RoSE provides highly interactive visualization, analysis, and storage of output from Red Storm and Sandia's capacity-computing clusters. Red RoSE and DSA-HPSS are integrated to provide archival storage of Red Storm output. The HPSS team and the RoSE Cluster team achieved a record-breaking performance in excess of 15 gigabytes per second on the RoSE Cluster's Lustre file system (Q4FY05).

Institutional file system capability demonstration: the ability to route the Lustre high-performance parallel file system between InfiniBand, Myrinet, and Ethernet networks was demonstrated. This capability is key to the deployment of the institutional file system to Sandia's capacity clusters, and the routed file system is now in production operation in the Feynman cluster.

The Data Discovery FCLib library was used to build metrics for analysis of W80 abnormal mechanical data, which were computationally tractable but more pertinent than the default metric of "volume of dead elements." These FCLib metrics were used to improve sensitivity analyses used in V&V reviews of analyst validation and QMU assessments.

The security team has deployed the Group Information Service, which is the final component necessary for meeting level 2 Milestone #1594. This work and the work done through the FOUS/Cyber Security Operations project, have completed all major deliverables, including a final report (Q3FY06). Additionally, the University of Michigan CITI Linux NFSV4 reference software, with some modifications provided by this project, was tested utilizing the new LDAP directory infrastructure.

R&D efforts in Remote Direct Memory Access (RDMA)-enabled networking and storage made significant strides this year. The OpenIB (now OpenFabrics to embrace InfiniBand and 10 Gbps ethernet RNICs) Linux software stack with our active participation is scheduled to announce its first official software release on June 13th. We successfully debugged and demonstrated cutting-edge RDMA-enabled storage technologies: iSCSI with RDMA enhancements (iSER), SCSI RDMA protocol (SRP), and Network File System (NFS)-over-RDMA (Q2FY06). In addition, we released a software emulated kernel iWARP and DAPL stacks to Source Forge to encourage early development of RDMA applications and developed a parallel file system model for OPNET for large-scale network simulation studies of the ASC environment (Q1FY06). We started the evaluation of the first 10 Gbps Ethernet NICs that support RDMA in hardware from NetEffect, a technology that promises to greatly increase the bandwidth per CPU cycle ratio of emerging 10 Gbps ethernet systems (Q3FY06).

The Data Discovery AVATAR project for terascale “commodity” pattern recognition and the scientific visualization project, ParaView, were both selected as projects to represent Sandia's Computer Science and Engineering program at the CSE External Review (Q2FY06).

## Facility Operations and User Support

### LLNL Accomplishments for Facility Operations and User Support

Major accomplishments during the fiscal year have been the completion of installation and integration of both Purple and BlueGene/L platforms. BlueGene/L is now in production and by the end of CY06 (1QFY07), Purple will be GA but started running significant applications in 2QFY06. Both of these efforts involved all parts of the subprogram—with facilities deployment, system administration, security plans, networking, and establishment of user support and documentation for the new systems. The major portion of each system is on the classified network, but there are also unclassified portions of each system. BlueGene/L is the world’s fastest supercomputer, according to the latest TOP500 list, released June 22, 2006. BlueGene/L performed 136.8 teraFLOPS, or trillion operations per second. Part of the BlueGene/L success was the implementation of a large federated gigabit Ethernet network needed by the Lustre file system.

Other significant accomplishments in FY06 include the retirement of the ASC White system and the installation of its replacement system, Rhea. Rhea is expected to be GA in 1QFY07. Also, the responsibility for LLNL classified email service was transferred to FOUS late in FY05 and the service was significantly upgraded in FY06.

### LANL Accomplishments for Facility Operations and User Support

The following systems operations accomplishments were noted in FOUS:

- Created 64 bit BPROC
- Converted several segments of Lightning/Bolt to 64 bit Linux OS; this transition should be complete by the end of FY06
- Integrated PaScalBB GPFS into secure network for access by Lightning/Bolt systems and future HPC systems
- Improved monitoring tools for HPC systems and currently working on implementing the LLNL SNMP-based monitoring tools to augment existing tools

The following network operations accomplishments were noted in FOUS:

- Upgraded X-Div/Nicholas Metropolis Center for Modeling and Simulation (SCC) Secure Network Backbone to 10GigE
- Upgraded DISCOM network to 2 x 10 GigE connections to Sandia/ABQ and LLNL and installed new 1 gigabit IP encryption devices (TACLANES)
- Upgraded links between HPC resources (Q, HPSS, Lightning, ViewMaster) to 4 port 10GigE trunks
- Installed new top-level network switch in to support higher bandwidth connections between Secure HPC resources and HPSS

The following network infrastructure accomplishments were noted in FOUS:

- Infrastructure and connectivity added to support segment upgrades to Lightning

- Infrastructure and connectivity added to support HPSS upgrades (new Linux Movers and Titanium tape drives)
- Infrastructure and connectivity added to support Panasas/PasScalBB upgrades
- Installed 10GigE and Infiniband equipment in network tested; network tested now has version of all production interconnects for HPC clusters at LANL

The following user services accomplishments were noted in FOCUS:

- The Application Readiness team was formed as a collaboration with CSSE and Integrated Codes to provide direct assistance to ASC code development teams in transitioning production codes to new platforms
- The Application Readiness team successfully established two major ASC projects in production mode on the Lightning/Bolt platform, paving the way for the transition of related workload from the Q machine to these systems

The following facilities infrastructure accomplishments were noted in FOCUS:

- The design for additional 2.4/4.8 MW of power at SCC was completed. This is in preparation for the actual power upgrade for the SCC that will be finished in FY07. The project was designed to add a total of 4.8 MW of machine-conditioned power to the SCC in two stages of 2.4 MW each. In FY07, the first 2.4-MW upgrade will be completed.
- Operations and maintenance continued for mission critical electrical and cooling equipment, including cooling towers, chillers, pumps, substations, and UPS that support the computer room equipment. One major addition is the assumption of system maintenance for the Lightning system by operating personnel. This includes running diagnostics and replacing parts of the Myrinet interconnect system on the Lightning computer.
- Other facilities project in FY06 included an engineering study for using water cooling in the SCC and Laboratory Data Communications Center (LDCC) computer facilities, an engineering study for the installation of an additional chiller in the LDCC, and the removal of Halon from all fire suppression systems in the computing facilities. The LDCC was completely converted to an unclassified computer room, with all classified computing consolidated to the SCC.

The following Roadrunner facility work accomplishments were noted in FOCUS:

- Made significant progress on facility work in preparation for Roadrunner. This included the design for additional 2.4/4.8 MW of Power at SCC. This is in preparation for the actual power upgrade for the SCC that will be finished in FY07. Other work in FY06 in this area included adding additional rotational UPS protection for the SCC and ordering equipment to build a second substation to support the SCC.

### **SNL Accomplishments for Facility Operations and User Support**

Completed successfully the Red Storm-related Limited Availability Milestones #17, #18, and #1313 (Q4FY05).

Completed ASC level 2 milestone #1594, *Deploy Security Infrastructure*, transitioning away from end of life DCE security environment (Q2FY06). Cyber Security Operations deployed an LDAP Directory services as the DCE registry replacement. This combined with the Group Information Service developed and deployed through the CSSE -IOSSN Security project completes the major deliverables for the milestone. The demonstration of the DCE replacement infrastructure has been completed, and the new Kerberos and

LDAP directory infrastructure is now being utilized by the ASC systems as well as Sandia's web-based applications.

Demonstrated early production status for Red Storm, serving critical calculations for LANL, LLNL, and SNL (Q1–Q3FY06).

Released production versions of Collaboration, Learning, Information and Knowledge (CLIK) tool, basic training for Red Storm (developers and analysts), tri-lab common path URL for High Performance Computing web presence, and expanded use of SARAPÉ web-based intersite account request process (Q4FY05, Q2FY06).

Completed the 7x Suite of benchmark runs on ASCI Red in preparation for the comparison on identical problem runs on Red Storm (Q4FY05).

Purchased and installed Software Quality / Development cabinets for Red Storm Risk Mitigation efforts supporting Red Storm acquisition and deployment for Limited Availability (Q4FY05). Completed addition of power to service fifth row of Red Storm and upgrade of memory, SeaStar replacement, and Dual Core processors (Q3FY06).

Initiated system monitoring and fault detection and notification activity in the Cyber Enterprise Management project covering all scientific computing systems.

Delivered production quality Red Storm cycles to several ASC milestone projects: LANL W76, Sandia V&V and W-76 Life Extension Program (LEP). Completed benchmark and scaling studies required for General Availability milestone (Q2FY06).

Completed installation and checkout of Thunderbird Cluster. Began Production use in Q1FY06. Linpack Benchmark results of 38.27 teraFLOPS ranked #5 in TOP500 in November 2005.

Placed a contract with Qwest to replace the AT&T point-to-point OC-48 (2.5 Gigabit per second) circuit that connected SNL /NM to LLNL. The new 10 Gigabit per second service completes the full tri-lab ring connecting SNL /NM to LANL and LLNL at four times the bandwidth for one-half the cost of the previous service. The new topology provides enhanced service capability (higher bandwidth) and reliability (failover).

The WAN team tested and qualified the new 1 Gigabit per second Type 1 Internet Protocol encryptors for operational deployment in the ASC WAN. The new encryptors are necessitated by moving from OC48 ATM to 10 Gigabits per second Ethernet technology. The team also completed the cutover from the AT&T point-to-point OC-48 link to temporary Qwest OC-48 service. This provided service continuity while Qwest executed the installation of the two 10-Gigabit-per-second links.

Carried out an ASC WAN IPv6 readiness survey in preparation to meet the OMB IPv6 mandate.

Cut over ASC production traffic from the temporary point-to-point OC-48 link to the new 10-Gigabit-per-second ASC ring network.

Transitioned ASC production traffic from ATM encryptors to the new 1-Gigabit-per-second IP encryptors. The removal of ATM services will facilitate the retirement of aging ATM equipment from the ASC WAN.

Published the IPv6 readiness report for the ASC WAN to meet the OMB mandate.

Completed an OPNET-based simulation model of the new ASC WAN ring network used to enhance management and operations. The model will allow us to examine potential network reconfigurations without impacting production.

## IV. Product Descriptions by the National Work Breakdown Structure

---

### ***WBS 1.5.4: Computational Systems and Software Environment***

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of NNSA. It strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms, which include capability, capacity, and advanced systems. Along with these powerful systems that ASC will maintain and continue to field, the supporting software infrastructure that CSSE is responsible for deploying on these platforms includes many critical components, from system software and tools, to Input/Output (I/O), storage and networking, to pre- and post-processing visualization and data analysis tools. Achieving this deployment objective requires sustained investment in applied research and development activities to create technologies that address ASC's unique mission-driven need for scalability, parallelism, performance, and reliability.

#### **WBS 1.5.4.1: Capability Systems**

This level 4 product provides capability production platforms and integrated planning for the overall system architecture commensurate with projected user workloads. The scope of this product includes strategic planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research for future systems.

##### **Capability Systems Deliverables for FY07**

- ASC Purple will be released for general availability for classified NNSA computing. This will complete the ASC Program's level 1 milestone #350.
- Livermore will begin plans in FY07 for a new petascale production capability system with an estimated delivery in 2010, coupled with a substantial initial delivery system late 2007 or early 2008. A component of this planning is to understand and document NNSA SSP requirements, which together with ASC requirements from all laboratories, are drivers for the performance and sizing specifications for an interim 2008 and final 2010 system.
- The Roadrunner Project will be managed through the Systems Requirements and Planning project and executed by means of several projects in this Implementation Plan as coordinated through the LANL Roadrunner Project Plan. The base system hardware acquired in FY06 will be integrated to provide 81 teraFLOPS of additional capacity computing resources. In FY07, the lab will augment the base system with advanced systems compute accelerators and evaluate this approach as a new

technology for petaFLOPS computing targeted to break the petaFLOPS barrier in FY08. Such heterogeneous computing architecture shows promise for providing a new technology option for future capability systems.

- Los Alamos will begin planning and preparations for the next generation of tri-lab capability computing.
- Sandia will complete the upgrade of all five rows of ASC Red Storm to 2.4 GHz dual-core Opterons. The system's memory will also be upgraded to a uniform 6 GB per processor (3 GB per core). The completion of this upgrade is a candidate level 2 milestone from Sandia for 1QFY07.

### **Strategy to Deliver Capability Systems for FY07–FY08**

- Partner closely with the vendors throughout the contract period, from design and build to acceptance and throughout the life of the system.
- Actively partner with off-site labs to develop tri-lab platform strategies and aid in the efforts at the other labs.
- Tri-lab ASC Platform Strategy team continues to work together with Federal ASC managers to determine system siting and other major considerations.

#### ***1.5.4.1.1 Purple (LLNL)***

The Purple contract, a collaboration of the tri-lab community and IBM led by LLNL, has delivered a 93 teraFLOPS capability to the SSP. It is the premier NNSA classified capability computer serving weapons scientists and engineers at all three NNSA labs. This system is a near five-fold increase in scale for a single compute system for NNSA. The major architectural components of Purple consist of a computational cluster, a high performance network infrastructure, and a cluster-wide storage subsystem. The heart of the ASC Purple cluster is IBM's pSeries POWER5 based SMPs. The high-performance network infrastructure is IBM's Federation 4-gigabyte switch. The cluster-wide storage subsystem is based on IBM's GPFS running over 2-gigabytes-per-second Fibre Channel. Access to the storage system from outside Purple will be accomplished using parallel file transfer protocol over 1- and 10-gigabytes-per-second Ethernet. Purple was delivered to LLNL in 2005, moved to the classified network, and fully integrated during 2006. In the second quarter of FY07, Purple will enter general availability—any user with a valid account can access the system. Purple is planned to have a five-year lifespan, with the likely end of life in FY10. System requirements and planning effort will start in the first quarter of FY07 for a new petascale production capability machine targeted for 2010 to eventually replace Purple, with a substantial initial delivery system in late 2007 or early 2008.

#### ***1.5.4.1.2 Systems Requirements and Planning (LANL)***

This project covers all aspects of program and procurement planning for future capability, capacity, and advanced systems and strategic planning for supporting infrastructure.

The main focus of the project is to define requirements and potential system architectures for future capability platforms that meet programmatic requirements and drivers. This project provides a focus for the various planning efforts and provides project management support for those efforts.

In addition to planning for future capability platforms, this project will focus in FY07 on the project management and procurement of the Roadrunner system to be deployed at Los Alamos.

The Roadrunner contract with IBM is scheduled to deliver near term capacity computing cycles and an option for significant advanced architecture systems in the next two years. The Roadrunner Project is based on acquiring, installing, and deploying over 81 teraFLOPS of a Base Capacity System to provide capacity computing cycles in the near term and options for acquiring and deploying an advanced architecture system that will provide up to a petaFLOPS of computing cycles to the weapons program. The advanced architecture hardware, along with the base system, will be combined to create a hybrid computing architecture that has the potential for significant improvements to the price/performance curve to help meet the computing requirements in the future.

The project consists of three phases. The first phase consists of acquiring and deploying the Base Capacity system (over 81 teraFLOPS in secure environment and 10 teraFLOPS in open computing environment) and will include delivery of early access advanced architecture technology for testing purposes. Sixteen connected units (CUs) will be delivered in FY06 and deployed in FY07 to provide capacity computing cycles to the nuclear weapons program. The major architectural components of the Phase 1 base system includes a computational cluster (X3755 Servers), a high-performance network infrastructure, and high speed access to a storage system using 10 GigE networks. Each CU consists of 144 nodes (4-socket dual-core Opteron systems) with 32 GB of memory per node. The CUs will be connected by a second stage of Infiniband switches to provide the high performance.

The Roadrunner Phase 1 system will be delivered in September and October of 2006 and will run the Acceptance testing before moving to the classified network in early calendar year 2007. The Roadrunner classified capacity system will be made available for selected applications. As part of Phase 1, the current generation of Cell Blade technology will be installed on one of the unclassified CUs for initial work on the advanced architecture system.

The second phase consists of deliveries for the assessment of the accelerator technology path leading to a go/no-go decision to exercise the third phase procurement of the full set of advanced architecture technology. These deliveries include software and technology upgrades to the original Cell Blades for further assessment. This decision will be made at Los Alamos following a review by the ASC program (to include tri-lab and other external participants). The formal assessment and decision for the Phase 3 system will be completed in early FY08. There are several criteria for this assessment. One of the most important is the potential applicability of this architecture to weapons applications. Specific algorithms will be analyzed on early delivery systems of the advanced architecture as described in the Roadrunner Advanced Algorithms and Assessments project (Volume 1).

The third phase is the option for the acquisition of the final Roadrunner system architecture that will add Cell Blade 2 (CB2) dual-Cell blades to the base system nodes. The advanced architecture part of the system is expected to provide up around 1.7 peak petaFLOPS of computing cycles. This phase consists of additional Base System Capacity (5 teraFLOPS) as well as all the advanced architecture hardware, delivered in stages, for completing the system. If this phase is exercised the final system acceptance test will be completed in FY08.

Currently, the status of the project is that the procurement process has been completed and has resulted in the selection of IBM as the system vendor for the Roadrunner system. The Roadrunner Statement of Work (SOW) is part of the subcontract and describes in detail the technical, delivery, and acceptance criteria for the Roadrunner system.

In FY07 the System Requirements and Planning project will provide overall project planning for the Roadrunner system. This will include system integration and management for the Phase 1 capacity system and oversight for the Phase 2 assessment in preparation for the decision to exercise the Phase 3 option. In the event that the Phase 3 option is exercised, this project will provide management of the platform procurement for the advanced systems accelerated system to achieve a “sustained” petaFLOPS Linpack benchmark. This project will provide overall planning and management of the Roadrunner project, which will be executed by means of various projects in this Implementation Plan, as coordinated through the Los Alamos Roadrunner Project Plan.

The Roadrunner procurement is the first NNSA large scientific computing platform procurement to use the DOE Order 413.3 project management process. Roadrunner was used to pilot the draft NNSA Office of the CIO Project Execution Model (PEM) for IT investments. Several critical decision points have been successfully achieved. This project will continue to prototype the application of the PEM process for the procurement and integration of Roadrunner system. We are extending the process to have a full Project Execution Plan for the entire project, which will be formally reviewed.

In addition in FY07, Los Alamos will continue planning for the petascale computing environment. Personnel from our technical staff have been working with other tri-lab networking and storage experts to address the petascale computing environment that will exist once petascale computational systems are available for tri-lab use in ASC. The initial effort is to identify unique requirements that go beyond the “just add more compute cycles, bandwidth and storage” approach and to identify any critical changes to the computing environment that may occur once petascale computing systems are deployed. In particular, it is important to understand and address the way computing, storage, and visualization will be done and the subsequent impact on technologies that will be needed to support these changes. To do this, it is imperative that the planners understand the current limitations of terascale computing and not only extrapolate those to the petascale realm but also look for different ways in which the overall environment may evolve. Planned activities in FY07 include continued collaboration with the tri-lab community in addition to local, industry, and university collaborations to develop a petascale network environment roadmap.

It is understood that from a network perspective, petascale computing platforms will require connectivity to other petascale resources such as storage, visualization, and archive systems. A network that scales up to the performance of these platforms is required to maintain a balanced, efficient system. The group will address new networking technologies such as 40- and 100-gigabit Ethernet, OpenFabrics/OpenIB InfiniBand initiatives such as RNIC (RDMA enabled NIC) software stack, and Open Message Passing Interface (MPI) messaging libraries with an emphasis on scalability issues. This planning effort will define the requirements, technologies, and architectures that will scale the network up to the petascale range. The network will be deployed to support the petascale systems before their implementation. Similar approaches will address the storage and visualization aspects of petascale computing.

The effort may also focus on enhancing distance computing and machine area network user availability through the development of automated, proactive, performance monitoring tools that encompass the ASC resources at each site and include the WAN environment.

#### ***1.5.4.1.4 Red Storm Capability Computing Platform (SNL)***

Red Storm is a tightly coupled massively parallel processor with a little over 40 teraFLOPS of peak processing capability. The machine uses AMD Opteron processors

and a custom, very high performance 3D mesh communication network. Red Storm has a total of nearly 10,880 Opteron micro-processors. Red Storm has over 30 terabytes of memory and 400 terabytes of high-performance local disk that is split equally between classified and unclassified use.

The Red Storm computer system was designed to provide for a relatively easy upgrade path. The Red Storm architecture provides for linear physical scaling up to 32 K nodes. Increasing the size of the machine can be done simply by adding additional cabinets and communication cables. The Opteron processors are high-volume commodity parts that can be replaced with socket compatible higher performance processors as simply as in an upgrade to a PC. The memory is high volume commodity DDR DIMMs that can be increased to 8 gigabytes per processor with currently available memory DIMMs.

Red Storm has several unique features among which are its Reliability, Availability, and Serviceability (RAS) system, its Red/Black switching capability, and its partitioned system software functionality. In effect, Red Storm has a separate parallel computer system to manage and monitor the main system. This system has its own network and processors and its own operating system. All major components (including RAS system components) in the system are monitored by the RAS system. All errors, recoverable and non-recoverable, are logged by the RAS system.

Red Storm's unique Red/Black switching capability makes it possible for the machine to be used as both a classified and unclassified capability computing resource. The machine has 10,368 compute nodes that have from 2.0 to 4.0 gigabytes of DDR memory and 512 Service and I/O nodes each with 6 gigabytes of DDR memory. (The compute nodes in the normally classified section of Red Storm have 4.0 gigabytes of memory, the compute nodes in the center section of Red Storm have 3.0 gigabytes of memory, and the compute nodes in the normally unclassified section of Red Storm have 2.0 gigabytes of memory.) The 512 service and I/O nodes are split into two separate 256 node service and I/O partitions, one for classified and one for unclassified use. Through the Red/Black switching capability the compute nodes can be either all classified or all unclassified, approximately 25 percent classified and approximately 75 percent unclassified, or approximately 75 percent classified and approximately 25 percent unclassified.

Partitioning of the Red Storm system software functionality provides a full LINUX operating system on the service and I/O nodes and a lightweight kernel operating system on the compute nodes. The lightweight kernel operating system provides a substantial performance advantage on the compute nodes while the full LINUX provides the full set of features users expect to see on logging in. The LINUX operating system used on the service and I/O nodes were enhanced to provide users and system administrators a single system view of the machine.

In the fourth quarter of FY06, Red Storm will be increased in size through the addition of a fifth row of cabinets. This will increase the number of nodes in the system by 25 percent. Also, in the fourth quarter of FY06 a processor upgrade to Red Storm will be started and completed in early FY07. The combination of an additional row of cabinets and the processor upgrade will increase the peak performance to about 125 teraFLOPS. Later in FY07, the memory on the compute nodes will be increased to 6 gigabytes per node (3 gigabytes per processor core). This will result in over 75 terabytes of memory in the system.

This project also includes the Red Storm maintenance contract.

Red Storm is a highly scalable computing system and during FY06 it has demonstrated the ability to provide excellent parallel performance on ASC Nuclear Weapons applications. Red Storm's large memory (over 30 terabytes), efficient lightweight kernel

operating system, and very high performance network have contributed to make it a very efficient large-scale capability computing platform.

During FY07, Red Storm will transition to General Availability and will continue to provide production capability computing cycles for the ASC Program. In FY07, an upgrade of the Red Storm system to approximately 125 teraFLOPS and 75 terabytes of memory will be completed.

Red Storm will be in General Availability and full production throughout FY08.

#### ***1.5.4.1.5 Next Generation Capability Computing Platform (SNL)***

Sandia's Red Storm follow-on project is intended to support the efforts required for a new tightly coupled massively parallel processor in time to take over Red Storm's capability computing demands when it is retired in 2010. This follow-on system will build on and improve the very successful priorities for the Red Storm system architecture, such as, it will use an improved, custom high-performance 3D mesh communication network, interconnecting up to 64 K commodity multi-core processors. This highly integrated massively parallel processor system will feature a dedicated RAS subsystem, a Red/Black switching capability, and partitioned system software. Partitioning of the system software functionality provides a full Linux operating system on the service and I/O nodes and a lightweight kernel operating system on the compute nodes. The Linux operating system used on the service and I/O nodes will be enhanced to provide users and system administrators a fault-tolerant single system image.

This project will leverage the solid technology foundation that was developed for Red Storm. While the Red Storm system is the most scalable and reliable ASC capability platform, there remains room for improvement. Specific areas to be addressed include improvement in the interconnect design to support much higher rates of message handling throughput. The move to multi-core commodity processors means that improvements in interconnect bandwidth and latency need to be much greater than a simple scale up of Red Storm's SeaStar interconnect performance. The parallel file system for a Red Storm follow-on system will also need a breakthrough design in order to remove file systems as the *Achilles Heel* of future capability platforms. Part of the solution could be addressed by the explicit integration of the Scalable I/O subsystem with the RAS subsystem.

In FY07, we will begin technical exchanges with potential industry collaborators for a Red Storm follow-on. We will also pursue technical exchanges with potential funding collaborators such as DOE/ASCR and NSF.

We expect to deliver a draft requirements specification document for a tightly integrated massively parallel processor system to be delivered in 2009.

Preliminary planned activities in FY08 include an RFP solicitation and contract award.

### **WBS 1.5.4.2: Capacity Systems**

This level 4 product provides capacity production platforms commensurate with projected user workloads. The scope of this product includes planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include the procurement and installation of capacity platforms.

### **Capacity Systems Deliverables for FY07**

- In FY07, if funding permits, the laboratories will continue to procure, install, and integrate capacity compute systems on the Scalable Unit (SU) model to meet programmatic requirements for resources into FY08 and beyond.
- In collaboration with 1.5.5.1.3, Sandia will integrate and deploy upgrades to Thunderbird when OpenIB and OpenMPI are ready.
- The Roadrunner base system hardware acquired in FY06 will be integrated in FY07 to provide 81 teraFLOPS of additional capacity computing resources. This is a component of the Roadrunner Project, which will be managed through the Systems Requirements and Planning project and deployed by means of several projects in this Implementation Plan as coordinated through the LANL Roadrunner Project Plan.

### **Strategy to Deliver Capacity Systems for FY07–FY08**

- Leverage the extensive experience fielding world class Linux clusters within the tri-lab community.
- Provide quick expansion with production-ready systems while meeting programmatic need for performance, capacity, and efficiency.
- Continue to support ASC tri-lab initiatives related to capacity computing, such as Tripod.
- Augment system integration capability with additional personnel and deploy stable hardware and software components for the Roadrunner capacity system. Enhancements to Roadrunner to achieve even greater computer power will also be evaluated in FY07 and FY08 with actual weapon simulation codes.

#### ***1.5.4.2.2 Capacity Integration (LANL)***

The Capacity System Integration project will continue system integration for all capacity systems at LANL. This effort includes completing the ongoing system integration of Lightning/Bolt for the weapons stockpile computing workload, providing support to the Roadrunner system integration, and to other ASC production and test bed systems. These efforts will address complex integration issues relating to high-end terascale computing environments and provide a direct link to future petascale compute environments.

#### ***1.5.4.2.3 Deployment of ASC Capacity Systems (SNL)***

The purpose of this SNL project is to quickly integrate and deploy additional ASC capacity resources that leverage Tripod and open source software tools and capabilities. Sandia's experience with Thunderbird will be useful for specifying capabilities and functionality that should be included in this potential procurement, integration, and deployment process. The Thunderbird experience is useful because any future ASC capacity systems procurement will likely use a large Infiniband interconnect.

The focus of this project is to identify the appropriate integration of Tripod system software with the broader open source software community tools and capabilities deployed on Sandia's commodity Linux cluster capacity systems. The priority will be on adopting a system software stack that minimizes the total cost of ownership. If necessary, Sandia is willing to trade scalability in its capacity systems for the cost savings associated with adopting a system software solution aligned to the mainstream Linux cluster community.

Planned activities in FY07 include participating in a tri-lab ASC capacity procurement process, if it is resurrected. We will also continue to participate in the Tripod software stack process, if it is continued.

The FY07 deliverable of this production is to stand up, integrate, and deploy capacity systems made of commodity compute nodes, commodity interconnect technologies, and commodity out-of-band management cards. Lab-developed capabilities (potentially based on the Tripod software stack) may be integrated with open source Linux cluster system software components to help commodity clusters scale up to larger problem sizes.

Preliminary planned activities in FY08 include retiring aging Linux clusters by replacing them with updated, multi-core cluster solutions.

### **WBS 1.5.4.3: Advanced Systems**

This level 4 product provides advanced architectures in response to programmatic, computing needs. The scope of this product includes strategic planning, research, development, procurement, testing, integration and deployment, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research, and the investigation of advanced architectural concepts and hardware (including node interconnects and machine area networks) via prototype development, deployment and testbed activities. Also included in this product are cost-effective computers designed to achieve extreme speeds in addressing specific, stockpile-relevant issues through development of enhanced performance codes especially suited to run on the systems.

#### **Advanced Systems Deliverables for FY07**

- Develop specific recommendations for how to improve the balance factors that will limit scalability of systems available in the 2012–2016 timeframe.
- Deliver the Roadrunner Phase 3 system, pending the decision to proceed with the procurement of advanced systems accelerator technology, to achieve a 1-petaFLOPS sustained benchmark measurement and for application to ASC weapons computing needs.

#### **Strategy to Deliver Advanced Systems for FY07–FY08**

- Partner with the National Security Agency (NSA), Office of Science (ANL), and IBM to leverage planning and funding activities.
- Pursue a collaboration with the NSA Advanced Computing Systems (ACS) program to address research and development efforts for advanced interconnects, memory performance enhancements, and advanced RAS capabilities.
- Explore production code re-factoring techniques that will allow the codes to make use of Roadrunner accelerator technology to improve performance.

#### **1.5.4.3.1 BlueGene (LLNL)**

BlueGene/L is a scalable architecture in which the computational power of the machine can be expanded through additional low-power, highly reliable building blocks without introduction of bottlenecks as the machine scales up. BlueGene/L offers a theoretical peak computational rate of 360 teraFLOPS through extreme scalability. With more than 65,536 dual processor nodes, 32 terabytes of memory, and a 900-terabyte Lustre file

system, BlueGene/L is the fastest supercomputer in the world with 280 teraFLOPS delivered on the Linpack benchmark, 207 teraFLOPS delivered on an advanced quantum molecular dynamics code, and in November 2005, it won every HPC Challenge award at the 2005 SuperComputing Conference.

In FY06, BlueGene/L demonstrated at least three science applications of interest to the SSP running at over 100 teraFLOPS sustained for days solving problems that were reported by the Secretary of Energy to Congress. In addition, in FY06 BlueGene/L was migrated to the classified network and at least two critical ASC applications were ported to the machine and delivered results to the program within two weeks. These applications delivered results that surpassed the same simulations on Purple in terms of number of processors used and the resolution of the problem solved. In addition, these results clearly demonstrated the benefit of predictive simulation to the SSP.

In FY07, LLNL will continue to work on BlueGene simulation support environments with the Office of Science. IBM will continue to provide ongoing hardware and software support for BlueGene/L in FY07.

#### ***1.5.4.3.2 Advanced Systems Procurement and Supporting Technology Research and Development (SNL)***

This project is focused on the research and development of specific capabilities to actively address the memory and interconnect imbalances that threaten to spiral out of control with the transition to 32-128 core commodity processors in the 2012–2016 timeframe. It is likely that the processor manufacturers (Intel, AMD, and IBM) will recognize on their own that for multi-core to be successful and effective, they will need to develop advanced memory subsystem capabilities to keep all their cores fed with data. While the processor community is adding new capabilities, we want to have a specific set of recommendations to provide support for remote load/store functionality and capabilities that support the integration of many thousands of individual processors into tightly integrated massively parallel processor systems. This latter set of recommendations could leverage Sandia's work in advanced networking capabilities to identify processor memory controller capabilities that support synergies with high performance interconnects.

In contrast to memory subsystem performance capabilities, the development of interconnect capabilities is not likely to be addressed by the processor manufacturers. To address this shortfall, Sandia will use this project to lead the development of an interconnect infrastructure for the 32-128 core commodity processors that will be available for the 2012–2016 generation of capability supercomputers. The performance goals are aggressive: 30 GB/s of bandwidth, 500 ns of MPI latency, and over 30 million messages per second of MPI message throughput. Message throughput is a network characteristic that industry has shown little concern for, but it is critical for getting "usable bandwidth" from the network when applications use short messages. While message throughput is the most aggressive of the goals, it is unlikely that industry will achieve any of these goals in this timeframe.

The proposed interconnect work is to create both the network interface (NIC) and router chips to reach these levels of network performance. The overall project will include specifying the architecture for both of these chips as well as designing the two ASICs (NIC and router) that will be required to implement the network. The final year will include the fabrication of both chips to be ready to provide to a vendor. It is expected that a vendor partner would be engaged to build a system around these network chips.

The goal of this project is to provide network capabilities that will enable the creation of petaFLOPS scale capability computers from commodity processors. It will deliver

30 GB/s of bandwidth per direction per processor socket with 500 ns of latency through MPI and MPI message throughputs in excess of 30 million messages per second. This will enable petaFLOPS scale capability systems to be built from well understood commodity processors at a scale that is well understood. This is in contrast with other approaches that propose higher risk specialized processors or extremely high degrees of parallelism.

FY07 activities will include specification of the NIC architecture and router requirements and development of the initial VHDL or Verilog code. This activity is considered critical for developing a useable multi-petaFLOPS supercomputer for the 2012–2016 timeframe, but early success may also impact the 2009 timeframe capability system.

Expected deliverables in FY07 include NIC and router specification in Verilog or VHDL.

Preliminary planned activities in FY08 include completion of all VHDL/Verilog code and beginning of verification. We expect 50 percent verification completion by the end of the year.

#### ***1.5.4.3.3 Blue Gene/P and BlueGene/Q Research and Development (LLNL)***

The NNSA partnership with IBM on advanced systems, led by LLNL, was expanded in FY06 to include the Office of Science (SC), led by Argonne National Laboratory. This collaboration between NNSA, SC, and IBM is currently working a research and development contract that targets the development and demonstration of hardware and software technology for 1-petaFLOPS and 10-petaFLOPS systems. The 1-petaFLOPS system is known as BlueGene/P and a technology demonstration will complete in FY07. The 10-petaFLOPS system is known as BlueGene/Q and this effort will make microprocessor choices in FY07. In FY08, this effort will make BlueGene/Q node architecture and interconnect architecture choices. The BlueGene/P hardware is based on an extension of the highly successful BlueGene/L architecture with faster nodes, more memory, faster interconnects and larger system scalability. The software approach to BlueGene/P is open-source collaborative development between IBM research, Linux Technology Center, E&TS division and ANL and the ASC tri-labs.

In FY06, joint NNSA, SC, and IBM technical working groups were formed to architect, design, and develop the software for BlueGene/P. In FY07, these groups will be producing code for test on the BlueGene/P hardware demonstration rack.

#### ***1.5.4.3.4 Roadrunner Phase 3 Procurement (LANL)***

If the Phase 3 Roadrunner contract option is exercised, Los Alamos will procure advanced architecture components (Cell Blades) to add to the base configuration to form the final Roadrunner system. One more CU will be added to the classified system. The projected peak performance of the Cell blades is 1.7 petaFLOPS. This will be accomplished by adding Cell Blade 2 (CB2) dual-Cell blades to 138 nodes of each of the 15 CUs in the classified environment. Several software deliveries are also scheduled that will provide the full system with a hybrid architecture using the IBM System X AMD Opteron-based host servers (the x3755) accelerated with IBM's Cell Broadband Engine ("Cell BE") blades.

### **WBS 1.5.4.4: System Software and Tools**

This level 4 product provides the system software infrastructure, including the supporting operating system environments and the integrated tools to enable the

development, optimization and efficient execution of application codes. The scope of this product includes planning, research, development, integration and initial deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include system-level software addressing optimal delivery of system resources to end-users, such as schedulers, custom device drivers, resource allocation, optimized kernels, system management tools, compilers, debuggers, performance tuning tools, run-time libraries, math libraries, component frameworks, other emerging programming paradigms of importance to scientific code development and application performance analysis.

In FY07 ASC will continue to collaborate with Argonne National Laboratory on the development of the message passing library MPICH2. This product won an R&D 100 award in FY05 and ANL will continue their work for support of ASC platforms, improvements to the library including work on performance, tool integration and I/O. ASC will also contribute to an interagency high-end software strategy assessment effort that is being led by the DOE Office of Science.

One industrial contract that this product area supports is the tri-lab Open | SpeedShop (O | SS) Path Forward effort. Building upon the delivery of version 1.0 by SGI at the end of FY06, this is a project to deliver an integrated suite of open source performance tools for the ASC Linux capacity systems. The architecture of the software is explicitly developed with expandability for future tools and experiments in mind. In addition, the design is such that it should be easily portable to other systems. In FY07, plans are being formulated to firmly establish this project as an open-source, open-development effort with broader community support, and we are investigating strategies to use Open | Speedshop on ASC capability platforms, such as Red Storm and Purple. The second industrial contract is with OpenWorks to extend the open source Valgrind memory tool to the major ASC platforms. There are several project deliverables, including support for Purple and BlueGene /L, and Red Storm (depending on licensing constraints) that will be completed in FY07.

### **System Software and Tools Deliverables for FY07**

- For the Purple Level 1 general availability milestone, this product will deliver final testing, verify, and support all aspects of the tri-lab user software environment. It will also provide all of the resource management software (LCRM and SLURM) for the system and work with users in their efforts to get scalable performance on the systems.
- Development and deployment plans for petaFLOPS systems, including Roadrunner, the Red Storm upgrade, and systems resulting from new RFPs will be delivered. Livermore is also assisting ANL and the Office of Science in software plans for their new BlueGene /P system.
- A new tri-lab resource management tool will be deployed.
- Environment software releases for a common capacity system user interface.
- Tools to help measure the Capability Performance Indicator (CPI) for Red Storm, continued development and support for the virtual node Catamount light weight kernel and runtime system software.
- Participate in the oversight of tri-lab contracts for tool development.
- Requirements and design documentation for a seamless user environment.
- Demonstration of Science Appliance, Open MPI, and Eclipse successfully inter-operating on a major application on a multi-architecture cluster.

- Initial development of tools to analyze and solve infrastructure issues.

### **Strategy to Deliver System Software and Tools for FY07–FY08**

- Focus on requirements from the ASC customers added together with the requirements that are derived from the platform strategy.
- Map the available resources against the requirements to deliver product results starts first with making the best possible use of the software provided and supported by platform partners.
- Provide the best tools and protect portability; there is a heavy focus on formal and de facto standards (for example, languages and communication libraries).
- Emphasize maturity and scale of the software for Purple and BlueGene/L, augmenting IBM software as needed. Move toward a software stack that provides more commonality between the tri-lab capacity systems, both from the user portability point of view and from the software development and support point of view.
- Apply lessons from the BlueGene/L system to software requirements for future petaFLOPS systems to influence development.
- Establish standardized metrics for measuring RAS performance and develop tools that allow RAS performance measurements on all ASC systems.
- Incremental development and integration of a seamless and productive user environment.

#### ***1.5.4.4.1 System Software Environment for Scalable Systems (LLNL)***

This project provides system software components for all the major platforms at Livermore, research and planning for new systems and future environments, and collaborations with external sources such as the platform partners, especially IBM and Linux vendors.

This project covers the system software components needed to augment Linux and the required proprietary operating systems that function in a manageable, secure, and scalable fashion needed for the Livermore platforms. Currently, this includes work on scalable system management tools, the security infrastructure to provide secure single-sign-on tri-lab access, and the resource management environment to queue and schedule code runs across the Livermore systems.

Primary activities in FY07 will focus on providing all capabilities of the project for the ASC Rhea cluster (White replacement) and preparation work for an initial development machine resulting from the petaFLOPS system procurement. Emphasis will be on completeness, manageability, and reliability of the system environments. Another major focus of activities in FY07 will be integration and deployment of the tri-lab Resource Management System on the ASC capacity cluster. An instantiation of the tri-lab RAS metrics implementation will be deployed. An identity management solution will be deployed within the Center, which is expected to create efficiencies in Livermore's account management services.

Preliminary planned activities in FY08 will focus on completing the system components for newly procured systems to the next level of maturity and scale. It is expected that there will be some activity around increased integration of the tri-lab RAS metrics implementation, along with increased emphasis on completeness of the open source environment. There will also be effort related to interaction with the ASC platform products to research and plan for future platforms.

#### ***1.5.4.4.2 Code Development and Performance Environment for Scalable Systems (LLNL)***

This project provides the code development tools for all the major platforms at LLNL, supports user and code productivity through the tools environment, provides research and planning for new tools and future systems, and provides collaborations with external sources of tools such as the platform partners, independent software vendors, and the open source community. Current partners include IBM, Etnus, ParaTools, Openworks, and the University of Tennessee. The project works directly with code developers to apply these tools to understand and improve code performance and correctness. These activities provide customer-based feedback needed to plan future improvements to the environment. The project also resolves bug and user trouble reports, including interactions with the software providers to fix problems.

The tools covered by this project include, but are not limited to, compilers, debuggers, performance assessment tools and interfaces, memory tools, interfaces to the parallel environment, code analysis tools, and associated run time library work, with explicit focus on the development environment for large-scale parallel platforms. This project does not cover the development environment for workstations or other special purpose systems.

Primary activities in FY07 will focus on providing all project capabilities for the ASC Purple system, the ASC BlueGene/L system, the ASC capacity environment, and preparations for new petaFLOPS systems. Emphasis will be on completeness and reliability of the code development and performance environment for Purple, including the full tri-lab environment required by the level 1 Purple General Availability milestone. Emphasis for BlueGene/L will be in providing a stable environment and learning how the environment/applications scale to the very large processor count. Feedback on the limitations of this environment will be used to provide input into development activities needed to make future petaFLOPS systems a reality. Emphasis for the capacity environment will be in providing the required common tri-lab environment. Emphasis for the petaFLOPS environment will be on developing open source partnerships with IBM and the SC on BlueGene/P efforts. This will include gathering application communication and system software requirements. Underlying all of these activities will be direct customer support for the ASC code teams as well as development of new techniques for improving the robustness and performance of ASC codes on all platforms, including new large-scale debugging techniques, mechanisms to discover performance trends automatically and automated source code analysis that improves performance or locates deprecated coding constructs. Performance improvement support for Alliance users on BlueGene/L will also be provided. User assistance, trouble reports, and interaction with vendors to improve or to fix software products will continue to be an important aspect of LLNL service to the users. Specific new tools to be deployed in the FY07 open source production environment include Open Speed Shop, the Open Trace Format, and Valgrind.

FY07 milestone deliverables for this project include certification of the Purple code development environment (part of ASC level 1 milestone 350), deployment of the tri-pod Applications Development Tools environment for capacity systems, and a plan for the petaFLOPS environments.

Preliminary planned LLNL activities in FY08 will focus on the code development and performance environment for new capacity and petaFLOPS systems. It is expected that there will be some activity around improving the commonality of the tri-lab code development environment software stack and increased emphasis on completeness of the open source environment. There will continue to be effort related to interaction with

the ASC platform products to research and to plan for future capability, capacity, and advanced architectures. General activities to improve existing systems and support users will continue.

#### ***1.5.4.4.3 System Software and Tools Development (LANL)***

This project produces production grade software for scalable Linux system management (Science Appliance), Scalable, Portable, Fault Tolerant MPI-2 Message Passing Library (Open MPI) with its supporting run-time environment (Open RTE), and a parallel aware IDE for enhanced productivity (Eclipse). This project covers the full software lifecycle of these products, from requirements gathering to software maintenance. Capabilities of the project include:

- Science Appliance: Provides a fast, reliable, scalable cluster management system. A new process management subsystem (Xcpu) is being developed that offers significant improvements over the existing bproc system. This work is being done jointly with Terrasoft, IBM, and others.
- OpenMPI: Provides a scalable, portable, high-performance, heterogeneous, fault-tolerant MPI-2 implementation. The implementation is highly configurable, using a component architecture approach to support a variety of implementations of a given feature within the same library. It is being developed as part of an international collaboration. This library provides support for messaging across multiple networks in a single job.
- Eclipse: Provides an Open-Source, multi-platform, industry-supported set of integrated developer tools for parallel computing, including an editor with advanced features such as code refactoring and MPI content assist, a parallel debugger, an integrated build system, and performance analysis tools. This work is being done jointly with IBM, University of Oregon, Monash University, and others.

Planned activities in FY07:

- Science Appliance: Ensure Xcpu works reliably and scales on LANL systems. Develop new tools for system monitoring and control.
- OpenMPI: Ensure LANL applications are able to successfully use OpenMPI for message passing and MPI/IO. The next release will include collective optimizations, scalability, and fault tolerance of the run-time system. While the release will be community open source, NNSA target machines will include Red Storm, LANL Myrinet based systems, and IB clusters.
- Eclipse: Ensure Fortran integration will support large, complex codes. Provide enhanced parallel debugger features and ensure debugger is scalable to large machines. Add remote services capability to allow Eclipse to be seamlessly integrated with the users' desktops.
- All tools: Software release and regression testing, cross functional test code development, and documentation. Second line software support, including bug fixes and full system debugging.

Preliminary planned activities in FY08 include:

- Science Appliance: Production integration of Xcpu, especially with a seamless user environment. Integrate new system monitoring tools and control tools into a production environment.
- OpenMPI: Integration of MPI and RTE libraries into a production environment.

- Eclipse: Integration of Eclipse-based debugging and analysis tools into a production environment and into the seamless user environment.

#### ***1.5.4.4.4 Seamless User Environment (LANL)***

This project will develop a user environment in which the user needs less knowledge of the detail of the physical make-up of the system. The user should be able to express the computational requirements and this environment will address those requirements transparently to the user, in order to make most effective use of the computing resources available.

Planned activities in FY07:

- Definition of what constitutes a seamless user environment
- Requirements document detailing what functionality the seamless user environment must provide
- Design document describing a proposed implementation for a seamless user environment

Preliminary planned activities in FY08 include beginning prototype implementations of a seamless user environment and ensuring they integrate with underlying software stacks and hardware.

#### ***1.5.4.4.5 Productive User Environment (LANL)***

This project will develop a framework to more efficiently develop and use ASC code products. The enhancements will address parallel I/O, run-time partition capability, and connection to visualization tools. We anticipate more efficient use of existing computer hardware, file, and storage systems. All the tools and source codes will be a part of ASC code products.

ASC code developers and DSW designers will be actively involved in the project. The output of the project will be measured through productivity increase of DSW designers for relevant calculations.

In FY07, we plan to develop a framework based on the above areas of interest. We will also perform measurements of productivity enhancements on relevant simulations.

#### ***1.5.4.4.6 Application Readiness Team (LANL)***

This project will deploy an applications readiness team that works directly with users to address application problems with the ASC computers. This team is needed to support users/code teams to do initial ports and deep problem determination of code issues on new platforms. This effort will facilitate the early use of the platform by production codes and avoid wasting resources through low utilization. The intent of the team is to bring increased systems/computer science expertise to bear and to supply computer scientists who know both applications and systems.

The tasks involved are:

- Drilling down on problems, taking full codes that are millions of lines of code in dozens of libraries, finding the offending few lines of code, creating reproducers, and working with systems experts to determine what can be done to fix the system or the code
- Taking hard application problems (hangs, slow downs, and segv's) and being able to drill down to an answer quickly

- Closing the loop with the regression testing so that issues found with the system can have the reproducer included in the regression tests for the system so the system regression tests get better over time through real problem recreation
- Closing the loop with the monitoring efforts to use the system issues found to make the monitors better
- Final integration testing of a new revision of a platform stack before returning to production use
- Getting on a new platform first to get applications up and running on that new platform, making it possible for much quicker successes on new platforms with little or no user frustration

The team will have access to the codes and will have an interface person in the code teams themselves who helps get answers about the codes and libraries.

In the past, this work (Q at Los Alamos) was done by the systems analysts from the vendor or a vendor's subcontractor. However, there is a need for permanent expertise at each laboratory for the long term as the activities are both valuable at new systems standup and also for ongoing maintenance, upgrade activity, and codes change over time as well. In FY07, the applications readiness team will both help to complete the port of production stockpile codes onto the Lightning/Bolt systems and will assist with transitioning the codes to the Roadrunner capacity system.

#### ***1.5.4.4.7 Software and Tools for Scalability and Reliability Performance (SNL)***

This project supports software development to address scalability and reliability performance on Sandia computational systems. This project has four components: 1) Red Storm system software development and support, 2) system software development for Red Storm follow on, 3) development of RAS performance capabilities, and 4) application performance benchmarking, analysis, and modeling for upgrade options and next generation systems.

The Red Storm system software development effort provided the technical foundation to meet the Red Storm Limited Availability level 2 milestones in 4QFY05. In FY07, we will provide system software support to upgrade Red Storm with dual-core Opterons. This project will also support the efforts to identify and resolve Red Storm I/O and parallel file system performance bottlenecks.

The RAS performance capabilities are focused on three areas:

- The development of standardized RAS metrics and tools to instrument current ASC capability and capacity systems (requires collaboration with and handoff to FOCUS 1.5.5.1.3). In FY05, a small project was started at Sandia to develop standardized definitions for RAS metrics that can be applied to both capability and capacity systems. A subset of such metrics has been implemented and is being reported on Sandia's Red Storm and Thunderbird systems. A key project goal has been to affect the broader HPC community towards standardized quantitative understanding of HPC RAS.
- Software tools that provide early fault detection and assist in automated fault recovery and diagnosis (also in collaboration with 1.5.5.1.3). In FY06, a CSRF project developed an active RAS tool, OVIS, for performing real-time monitoring and analysis on computational clusters. This statistical approach to cluster monitoring and analysis can enable earlier detection of problems than the traditional threshold methodology.

- RAS subsystem design specification to prepare for the Red Storm follow-on procurement, 1.5.4.1.4. This activity will also support the development of formal specifications and design for a RAS subsystem to help ensure the 2009 capability system has better RAS functionality and performance than Red Storm.

The application performance benchmarking, analysis, and modeling portion of this project will provide the ability to assess component and sub-system tradeoff for future system architectures under consideration by Sandia, and will allow the determination of the performance affects caused by system changes. The measurement of application performance will use tools developed and supported by 1.5.4.4.4. The use of application performance and benchmarking to improve the performance on existing system like Red Storm is a natural area for collaboration with FOUS 1.5.5.2.3 User Support Services product.

This project will provide the core support for Red Storm's Catamount Virtual Node lightweight kernel and runtime software. It will also support the integration of RAS capabilities into the scalable I/O and Lustre parallel file system to improve the performance and reliability of Red Storm under production application I/O loads. Tools are needed to measure our standardized RAS performance metrics in a manner that allows comparison across different systems.

Planned activities in FY07 include supporting the upgrade of Red Storm to dual-core Opterons with Catamount Virtual Node; continuing to address system software performance and reliability issues; and developing application performance models that quantify the benefit of system enhancements. SNL will continue to grow the initial effort to establish standardized metrics for measuring RAS performance and develop tools that allow RAS performance measurements on all our ASC systems. We are working to define a tri-lab Q4FY07 level 3 milestone to produce a RAS specification and metrics implementation useful for streamlining procurement and facilitating data-driven decision making. Existing tools will be enhanced to incorporate low-level component and environmental data combined with information from error logs and bug trackers to identify correlations to enable fault diagnosis and to provide troubleshooting decision support for system administrators (in collaboration with 1.5.5.1.3).

Expected deliverables in FY07 include system and runtime software to address the technical issues of paramount concern to ASC and running ASC applications on our capability and capacity systems with sufficient scalability, reliability, and application performance to meet our SSP requirements. The Sandia-specific deliverables for this product will include tools to measure our RAS performance.

Preliminary planned activities in FY08 include the extension of Catamount to Quad-core Opteron support and socializing our RAS metrics beyond the tri-lab ASC community.

#### ***1.5.4.4.8 ASC Programming Tools Project (SNL)***

This project is concerned with providing programmers tools to aid them in making their code reliable and efficient. Programming tools for the ASC program have unique requirements including coexistence with MPI, scalability to tens of thousands of processors, and operability on ASC platforms, which, in some cases, run operating systems unique to high performance computing platforms.

The tools required fall broadly into two categories, performance analysis tools and debugging tools. Performance analysis tools can be separated into tools for examining how well a program utilizes a processor and tools for examining how efficiently a program uses the communication network. Debugging tools can be divided into traditional debuggers, memory debugging tools, and code coverage tools. Traditional

debuggers (for example, TotalView) permit programmers to step through the execution of a program and inspect data. Memory debugging tools (for example, Valgrind) provide more detailed information than traditional debuggers by catching illegal memory reads and writes earlier in program execution than traditional debuggers can. Code coverage tools (for example, Javelina) make sure all code is executed during regression testing to ensure all potential bugs are exposed.

FY07 activities include a complete suite of application performance analysis and debugging tools for Sandia application development environments, by supplementing vendor-supplied tools and developing new functionality where needed. The SNL ASC Programming Tools Project is involved in the following activities to ensure ASC programmers have the necessary tools to provide efficient and reliable programs:

- Participate in the oversight of tri-lab contracts for tool development. Contracts that are currently in place include OpenWorks (Valgrind) and SGI (Open | Speedshop).
- Explore better ways to characterize application performance by developing performance diagnostics that guide programmers quickly to the source of performance problems. This work will leverage both the Sandia-developed VProf tool and the Open | SpeedShop framework, as well as existing tools such as mpiP.
- Selectively deploy preliminary tool releases to code teams for evaluation and testing. This work will be done primarily on Thunderbird and Red Storm and information will be made available to developers through the Tools Portal:  
<http://tools.ca.sandia.gov>.
- Continually familiarize ourselves with existing and new tools, developing training materials and usage guidelines for the tools. This will be particularly important for Open | SpeedShop and Valgrind, which will provide significant new capabilities in FY07.

This work will be done in close collaboration with Los Alamos and Livermore to ensure we can efficiently provide programming tools with little redundant effort.

FY07 expected deliverables include:

- Deployment of O | SS. Thunderbird deployment should be straightforward. Red Storm deployment is contingent on funding of the Phase II O | SS and on Cray fixing PAPI bugs on Red Storm.
- Deployment of latest Valgrind releases. Thunderbird deployment should be straightforward. Red Storm deployment is contingent on finding additional funding for completing Catamount support in Valgrind and an export license being granted for Catamount.
- Performance tuning workshop for ASC code developers.
- Preliminary tools for enhanced performance diagnostics.

FY08 preliminary planned activities include reviewing successes and failures of the latest tool features based on users' experiences during the workshop and planning work to correct problems. We will also deploy tools with enhanced performance diagnostics for the ASC code teams and continue with contract oversight and training activities.

#### ***1.5.4.4.9 Software for Heterogeneous Systems (LANL)***

The Los Alamos Roadrunner system in its final configuration will be a hybrid heterogeneous architecture of IBM Cell Blades attached to typical AMD Opteron nodes which are interconnected with InfiniBand to form a traditional cluster. This system thus

includes three distinct processing units and three memory subsystems that must be programmed and intercommunicate. These three levels include: the Opteron CPU, the Cell's PPC CPU and the Cell's eight SPE vector units, and the Opteron's memory, the Cell's common memory, and the small local store memory of each SPE. There is no single language, compiler, or tool that can treat these all within a single program source code.

IBM currently provides a primitive capability for programming a single standalone Cell chip, utilizing the PPC and SPE's and sharing a common memory subsystem. This requires explicitly coding two interacting programs and using Cell intrinsic functions. On the other hand, the necessary programming models and system and runtime software for having the Cell Blades compute and intercommunicate with the Opterons does not yet exist at all and is being developed by IBM under the Roadrunner contract. This project is working with IBM to architect, develop, test, and evaluate those tools and runtime libraries, with IBM as the implementer of the actual libraries.

This project also includes an important effort to model the new parallel hybrid Linpack code IBM is developing to achieve the 1-petaFLOPS sustained performance target. This sub-project is tasked with developing a model for this new hybrid Linpack and using it to extrapolate full-Roadrunner performance and for providing feedback to IBM about limitations in the new code so that adjustments and improvements can be made.

Lastly, this project includes a long-term component to plan and develop programming tools to improve beyond the near-term brute-force programming described earlier. There is the need for a development environment that can analyze, identify, and possibly refactor code segments for Cell acceleration, and to provide debugging and performance data collection of hybrid Roadrunner codes.

In FY07 the project will:

- Define a programming API for Opteron to Cell Blade hybrid computation and communication and write a functional specification document.
- Track, test, and evaluate the IBM developed Opteron to Cell Blade runtime libraries on the early access AAIS Roadrunner cluster.
- Track, test, and evaluate the IBM provided Cell Blade programming environment.
- Develop and apply a performance model for the new IBM hybrid Linpack.
- Start developing prototype tools and plans for next generation programming environments.
- Document and present results in support of the Roadrunner Phase 3 "Go-No-Go" assessment review(s).

Activities in FY08 will include:

- Present at the Roadrunner Phase 3 "Go-No-Go" assessment review(s).
- Continue Linpack modeling efforts to include the first wide-scale integration of Cell Blades in the final Roadrunner configuration and scale up to the full system 1-petaFLOPS Linpack run.

Focus the long-term programming environment efforts at strategic areas that need attention as demonstrated by the use and testing of the early brute-force hybrid programming tools.

#### **1.5.4.4.10 Application and System Performance Analysis (LANL)**

This project concerns itself with the performance of large-scale systems and applications. The application workload consists of realistic, full-blown applications—both classified and unclassified—from the tri-lab community, with an emphasis on applications of interest to LANL. The systems under consideration are all platforms of interest to the ASC program, including capacity systems, capability systems, and advanced architectures. The goals of the project are to provide LANL, the tri-lab community, and DOE with a deep understanding of the complex relationship of application characteristics, architectural features, and performance in contexts of prime relevance to the ASC program.

Planned activities in FY2007 include the following:

- Performance and scalability analysis of Roadrunner. Roadrunner, Los Alamos' latest capacity-system procurement, will go online with the base system in FY07. As part of Roadrunner's acceptance testing, we will benchmark the system at each stage of installation (unclassified half system, classified half system, classified full system) and analyze if the system's performance is in line with that promised by the vendor.
- Optimization of InfiniBand process mapping. An increasing number of clusters across the weapons complex use InfiniBand as the interconnection network. We plan to investigate the impact on performance of application-centric process-to-processor mappings in the context of InfiniBand clusters. Network optimization will be performed, based on improved routing schemes.
- Application workload analysis. We have a suite of accurate, validated performance models for various applications of interest to Los Alamos and the tri-lab community. Our intention is to expand this suite to include a more robust model of Los Alamos' Milagro application (an initial prototype model having been developed for FY06) and a classified LANL application.
- Performance prediction for the Red Storm upgrade. Sandia is investigating upgrading the processors in their Red Storm system from single-core Opterons to dual-core Opterons. Using our application-centric performance analysis, we plan on informing Sandia and NNSA how much performance improvement various ASC applications can be expected to see from the processor upgrade.
- Tools for large-memory jobs. Los Alamos' applications have large per-process memory footprints. We intend to begin a development effort in FY07 to investigate the benefits of a more flexible allocation of cluster memory to parallel applications.

#### **WBS 1.5.4.5: Input/Output, Storage Systems, and Networking**

This level 4 product provides I/O (input/output, or data transfer) storage infrastructure in balance with all platforms and consistent with integrated system architecture plans. The procurement of all supporting subsystems, and data transfer, storage systems and infrastructures occurs through this product. The scope of this product includes planning, research, development, procurement, hardware maintenance, integration and deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include high-performance parallel file systems, hierarchical storage management systems, storage-area-networks, network-attached storage (NAS), and HPSS or future hierarchical storage management system disks, tape, robotics, servers, and media. This product also includes relevant prototype deployment and testbed activities. Projects and technologies in the advanced networking and interconnect areas shall include networking and interconnect

architectures, emerging networking hardware technologies and communication protocols, network performance/security monitoring/analysis tools, and high performance encryption and security technologies.

### **I/O, Storage Systems, and Networking Deliverables for FY07**

- Provide appropriate throughput to/from storage and storage capacity, appropriate access to and capacity for global file systems, and appropriate inter-site access to resources via encrypted WANs.
- Provide performance-optimized and production-hardened I/O, storage, and local and WAN for the Red Storm computational, visualization, and data services resources. The file systems project will assist application and application library developers with performance optimization for the parallel file system and I/O libraries on Red Storm.
- Develop, through requirements, design, and code phases, HPSS Release 7.1 and actively perform the code review process on this software. HPSS R7.1 is focused on file create and transactional performance in support of petaFLOPS environments.
- Deliver a completed 25-gigabyte-per-second infrastructure for the classified Red Storm to the data services cluster; demonstrated Red Storm access at the committed rate of 100 megabytes per second from tri-labs via the DisCom network; a robust, integrated, tri-lab security infrastructure using Kerberos and LDAP services; an OpenIB open source InfiniBand software stack that meets ASC HPC performance requirements; and a number of analyses and reports on next-generation parallel file system and networking technologies for a petaFLOPS computing environment.
- Continue work with Red Storm and Lustre vendors to improve Lustre performance on Red Storm; assist application developers and users with I/O optimization for Red Storm and Sandia applications on Purple; continue research and development in collaboration with industry and academia in the areas of parallel file systems, I/O libraries, POSIX standards for parallel I/O, RDMA, data transport protocols, and networking technologies to prepare for petascale architectures.
- Deploy HPSS version 6 with full Linux support and no DCE dependencies; develop HPSS version 7 under the auspices of the ORNL HPSS contract with an emphasis on small file performance improvements; and continue development support for the transfer agent, PFTP, and LDAP integration.
- Research areas to address expected deficiencies and analysis of petascale networking environment technology roadmap; further development of OpenFabrics/OpenIB InfiniBand and Open MPI software to support dynamic and quasi-dynamic routing, subnet management scalability, topology awareness, and multicast enhancements for MPI collectives; and enhanced problem diagnostics and performance monitoring tools.
- ORNL HPSS developers, operating under an ASC Field Work Proposal (FWP), will support HPSS development tasks as directed by tri-lab ASC HPSS technical committee members at the level of 1.5 FTEs. This effort will support: 1) storage system management development and maintenance, 2) HPSS core server development, and 3) Level 2 support of ASC sites, primarily in the areas of the SSM and the core server.
- Improvements to the Open Fabrics Alliance InfiniBand software stack to improve performance, functionality, and stability will include better tools for performing diagnostics and management, support of higher communication rates, and quality of service. The tri-lab OpenIB/OpenFabrics contract will deliver scalable management

and subnet management/administration tools and capabilities, integrated into the official release.

- Deploy additional file transfer agents to meet transfer bandwidth requirements.
- Deploy v6.2 of HPSS in the secure network.
- Document and test v7.1 of HPSS.
- Deploy parallel storage interface with HTAR feature.
- Deploy OpenIB on new InfiniBand clusters.
- Implementation of the next phases of the PaScalBB.

### **Strategy to Deliver I/O, Storage Systems, and Networking for FY07–FY08**

- Ensure that research efforts and architecture and procurement decisions for these subsystems are made with a complete understanding of vendors' R&D and product roadmap and pricing; a thorough characterization of the product performance, reliability, manageability, and interoperability through testing and analysis; sustaining key industrial and academic collaborations; and a balanced and integrated planning effort considering the requirements all resources and funding allocation.
- Work in close collaboration with commercial vendors and the HPC community to influence technology maturation and leverage the investments of a very large community.
- The IOSSN project works with industry and academia to deliver scalable I/O infrastructure for current and future petascale architectures, in the form of parallel file systems, I/O libraries, RDMA, data transport protocols, and POSIX standards for parallel I/O; the project will continue development of HPSS with specific focus on integration into the Sandia ASC environment; and it will focus on networking technologies for the next-generation ASC petascale environment that must support a diverse set of I/O and interprocess communication requirements.
- Continue to address medium- and long-term user data storage requirements through a balanced mix of leading edge systems and versioned software development.
- Procure and install major file and archival system components in support of Roadrunner, to support base system in FY07 and accelerated system in FY08.

#### ***1.5.4.5.1 Archival Storage (LLNL)***

The Archival Storage Project provides end-to-end long-term, reliable, high performance, archival storage services to ASC customers. This includes all development, deployment, and support of scalable archival storage software and interfaces for tri-lab ASC customers on the unclassified and classified networks. It also includes the selection, procurement, deployment, support, and maintenance of all archival storage hardware and storage media and the ongoing technology refresh and data stewardship necessary to safeguard and maintain the fruits of ASC platform and science investment.

The Archival Storage Project provides state-of-the-art archival storage of data for ASC customers at all three laboratories. The project develops, deploys, and maintains the HPSS in concert with five DOE laboratories and IBM Global Services. HPSS provides a unique blend of scalable, parallel archival storage interfaces and services to customers running at all three ASC centers. Rather than constraining data transfer to the speed of a single storage device, HPSS is designed to distribute data across a configurable amount

of storage units and to remove other limits to scaling including number of files, directories, and concurrent users.

To supply the performance necessary to offload ASC platforms and not hinder computation, a world-class array of storage hardware is deployed and supported underneath HPSS. This includes high performance disk arrays, tape subsystems, mover nodes, storage-area-networks, networks, robotics and petabytes of media. Together, this hardware and software supports unlimited storage for an unlimited amount of time at speeds in excess of 2.5 gigabyte per second.

In FY07, the development arm of the project will focus on engineering HPSS Release 7.1 (R7.1), which will focus on greatly improving file create performance and transaction rates. Operationally, HPSS Release 6.2.1 (R6.2.1) will be deployed on both unclassified and classified networks through two carefully planned rollout efforts. The hardware and software required to support and expansion of ASC platforms will be deployed and supported, including multiple clusters of Linux mover nodes front-ending T10K and LTO-3 tape drives. The local HPSS team will work to integrate HPSS with the centers' scalable global file system.

In FY08, the development of HPSS R7.1 will be completed and the product will be released. While fielding the required hardware in support of ASC platforms, HPSS R7.1 will be deployed in both LLNL HPC centers. The tri-lab development process for HPSS R8.1 will be launched. This release will focus on the dynamic scaling core servers for a petascale environment.

#### ***1.5.4.5.2 File Systems (LLNL)***

The File Systems project provides for the development, testing (feature, capability, performance, and acceptance), procurement, integration, and ongoing support of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included is the continuing development and support of Lustre as a fully featured file system for the range of ASC capability and capacity platforms, the deployment and support of GPFS on the ASC IBM platforms, and the deployment and support of ubiquitous NAS services for home, project, and scratch space. A critical component addressed by the Scalable I/O effort (SIOP) is the testing, benchmarking, and support of various programming interfaces for parallel I/O, as well as collaboration with the tri-labs and academic research partners to address future ASC high-performance I/O needs.

This project deploys and supports large-scale, high-performance Lustre and GPFS file systems for ASC platforms as well as high-availability NAS file systems for home and project space, and scratch space for serial capacity clusters. Total disk capacity at the end of FY06 under Lustre is approaching 3 petabytes with the addition of BlueGene/L and an additional 2.8 petabytes under GPFS with Purple. NAS services provides highly-available home and project space using the NFSv3 protocol that is shared among all resources on each of the classified and unclassified networks to enable transparent file sharing between platforms. SIOP addresses the problem of I/O capability by supporting several standard programming interfaces for parallel I/O and by helping developers of applications and high-level I/O libraries use parallel I/O effectively as well as providing expertise in I/O and file system stability and reliability testing, and performance characterization.

The planned activities in FY07 are to deploy into production Lustre version 1.6 providing Kerberos support and other features. We will begin deploying NFSv4 as a replacement for current NFSv3-based NAS services; continue ongoing support of currently deployed NAS services and capabilities; and investigate and analyze the

expansion of existing environments, based on center-wide increases in compute capabilities. We will also investigate and analyze new technologies in the areas of NAS data redundancy (on center critical systems) and NAS backup technologies, as well as parallel I/O and file system research and advancements. In addition, we will maintain parallel file system and I/O library support for users.

Preliminary planned activities in FY08 include deploying into production Lustre version 2.0 providing clustered metadata, support for a larger number of objects and clients, and enhanced fault tolerance. We will provide investigation and analysis in expanding existing environments, based on center-wide increases in compute capabilities such as BlueGene/Q.

#### ***1.5.4.5.3 Networking and Testbeds (LLNL)***

The LLNL Networking and Testbeds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect sub-systems in support of current and future systems and environments. This work relies heavily on an adequately provisioned testbed, skilled staff, and collaborations with vendors.

This project will test various hardware and software components to quantify the features, performance, reliability, security, and interoperability of the products and broader technology base. The information acquired as a result of this project will be used to help determine an integrated architecture and resultant procurements for these sub-systems.

In FY07, in support of planning for future petaFLOPS systems, this project will perform research and testing for technologies and products pertaining to interconnects, local and WANs and IP NSA Type 1 encryptors, and file system servers, clients and disks. Of particular interest in FY07 are cost-effective, large 10GigE switches, additional features in InfiniBand, other interconnect and 10GigE related technologies, and 10GigE NSA Type 1 encryptors. Results of these efforts will optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems.

Preliminary planned LLNL activities in FY08 will further leverage tri-lab activities in I/O related hardware and software, and seek to improve the reliability, performance and manageability of the I/O sub-systems in production. Additionally, this project will perform research and testing to determine what technologies and products should be considered for insertion into production to meet the growing I/O performance and capacity requirements.

#### ***1.5.4.5.4 Archival Storage Services (LANL)***

This project includes all services for archival storage operated by Los Alamos for the purpose of supporting ASC customers from Los Alamos, Livermore, and Sandia. These services include HPSS system configuration, data security, system administration, system operation and monitoring, problem determination and error checking, database administration, and LDAP design/implementation. The project works with users to troubleshoot problems experienced with storing or retrieving data from the archive and also works with users on implementing need to know and other data security issues. This project is also required to collaborate with archive storage software developers for supporting new tape and disk technologies, meta-data issues, first-level support for corrective action, and with network engineers for implementing a data path from super computers/clusters to the archive that meets ASC bandwidth targets. Thrust areas include:

- Archive system support (on-going): Conduct daily system administration and continual monitoring of the archive database and HPSS meta-data server, tape and disk servers, and robotic tape libraries. Install operating system patches, driver upgrades for new tape and disk hardware, network configuration, disk and tape management, and HPSS configuration. Monitor systems for security events. Analyze and improve performance of key subsystems, including data movers, network switches, and meta-data processing. Develop scripts and tools to aid in administration and operation. Design and deploy new infrastructure architectures and services. Perform technology refreshes, which require the movement of data to newer tape media. Maintain security plans and tests for existing archive storage infrastructure components and modify plans as required for new production resources.
- RAS (on-going): Continuously improve the end-to-end level of service as seen by the users. Conduct ongoing studies and improvement projects for decreasing the time it takes for users to archive data from super computers. Implement new technologies that aid in decreasing the total time it takes users to archive their data. Improve diagnostic and monitoring capabilities. Provide on-call support 24/7 for corrective action for HPSS and PSI issues.

Planned activities in FY07 include:

- Deploy additional file transfer agents to meet transfer bandwidth requirements
- Deploy v6.2 of HPSS in the secure network
- Document and test v7.1 of HPSS
- Deploy PSI with HTAR feature
- Upgrade disk technology
- Test and deploy titanium tape drives
- Develop and implement formal and informal customer satisfaction metrics and measurement techniques
- Enhance RAS metric collection, reporting, analysis and management

Preliminary planned activities in FY08 include:

- Continued support of the production archive
- Identification and implementation of clustered data movers
- Upgrade HPSS core server
- Document and test v8.x of HPSS
- Test second generation titanium tape drives
- Implement improved RAS reporting

#### ***1.5.4.5.5 HPSS and PSI Development (LANL)***

This project includes services for HPSS and PSI software development by Los Alamos for the purpose of supporting ASC customers from Los Alamos, Livermore, and Sandia. These services include collecting user requirements for changes and upgrades to HPSS and PSI, developing plans for implementing user requirements into the codes performing the design and development work for upgrading the codes, and providing second-level support for the archive storage deployment team. The project works with the consulting office and archive storage deployment team to troubleshoot problems

experienced with storing and retrieving data from the archive. The HPSS portion collaborates with tri-lab developers for implementing solutions that meet ASC requirements for all three labs. The PSI portion collaborates with tri-lab colleagues on user interface issues potentially providing a common user interface for tri-lab archive storage access. Thrust areas include:

- Deliver and help with deployment of HPSS v7.1
- Provide new version of PSI with HTAR support
- Collect requirements for next releases of HPSS and PSI
- Create plans for implementing user requirements
- Design, develop, and test next versions of HPSS and PSI
- Participate in meetings for tri-lab, common user interface to HPSS
- Work with customers in developing archive storage performance metrics

Planned activities in FY07 include:

- Design, develop, and test HPSS release 7.1
  - Small file performance improvements
  - Multiple TCP streams for single stripe in each transfer
  - Mover affinity
- Provide second level support for archive storage deployment team
- Provide HPSS user, installation, and administration manual updates
- HPSS v8.x requirements and design document
- Provide new version of PSI with HTAR capability
- Make decision on tri-lab, common user interface
- Improve time to last byte for code projects A and B
- Evaluate file lifetimes for expiration of data within the archive
- Develop and implement formal and informal customer satisfaction metrics and measurement techniques

Preliminary planned activities in FY08 include:

- Continued level-2 support for the production archive
- HPSS 8.x development and testing
- Initiate next generation archive planning
- Continued development and support of PSI
- Consider making ACL's portable and seamless across the tri-labs
- Design, develop, and implement HPSS/PSI resource manager
- Support for resource manager supported by PSI
- Implement improved metrics reporting

### ***1.5.4.5.6 Scalable Input/Output and Storage Systems (SNL)***

The SNL Scalable I/O and Storage Systems project provides scalable I/O infrastructure in the form of highly parallel file systems and I/O libraries for Red Storm and large-scale Linux clusters, high performance HPSS archival systems, high performance data transport protocols and technologies (for example, iSCSI block transport over RDMA), and security R&D on these components to ensure proper authorization and authentication to ASC data.

The Scalable I/O (SIO) team works with vendors and researchers to provide reliable, high-performance, easily used scalable I/O libraries and file systems that make optimum use of disk I/O rates of 100–200 gigabytes per second and beyond. It also educates users on how to ensure optimum I/O performance from their applications. The library R&D contributes to the integration of current and future compute platforms, visualization rendering engines, and data management pre- and post-processor servers. The project works to ensure that higher-level I/O libraries effectively use lower level libraries; in particular, HDF5, MPI-IO, and industry and research file systems must work well together on all ASC platforms.

HPSS is a hierarchical storage management system comprised of disk cache and tape storage systems that provided end-to-end, long-term, reliable, high-performance archival storage services to ASC customers. The project develops and deploys HPSS in concert with five DOE laboratories and IBM Global Services. Sandia expertise supports the Red Storm Transfer Agent code, parallel FTP (PFTP) code, and LDAP integration.

Data transport protocol and technology R&D is focused on RDMA and offload-enabled technologies that are used to implement scalable, high-performance, distributed storage systems over high-speed networks such as InfiniBand (10+ gigabits/sec) and 1/10 gigabit Ethernet. Both emerging and research prototype block level (iSCSI, iSER, SRP) and remote file system (parallel NFSv4, NFSv4-over-RDMA) transports are being investigated in a research testbed.

Cyber-security analysis is investigating underlying security mechanisms for storage systems in the ASC environment. As storage systems continue to transition from being controlled peripherals of single computer systems to highly distributed, intelligent, network-facing systems in their own right, the challenge of providing an effective and high-performance security services becomes even more critical for the ASC program.

Planned activities in FY07 include analysis R&D of the NFSv4 security model for parallel and RDMA-enabled data access extensions in the ASC wide-area network environment; performance and scalability testing of Lustre and I/O libraries (parallel HDF5 and eXtensible Data Model and Format (XDMF) on Red Storm; assisting application and library developers with performance optimization of the parallel file systems and I/O libraries for Red Storm and SNL applications on Purple; continued R&D with industry and academia on parallel file systems, I/O libraries (MPI-IO, HDF5), RDMA and offload-enabled data transport protocols (pNFS, NFS-over-RDMA), POSIX standards for parallel I/O, PVFS2, iSCSI, iSER, SRP, and DAPL; and continued development of HPSS with specific focus on integration into the Sandia ASC environment.

Major FY07 deliverables include performance and scalability testing of Lustre and I/O libraries (parallel HDF5 and XDMF) on Red Storm; assisting application and library developers with performance optimization of the parallel file systems and I/O libraries for Red Storm and Sandia applications on Purple; continued R&D with industry and academia on parallel file systems, I/O libraries (MPI-IO, HDF5), RDMA and offload-enabled data transport protocols (pNFS, NFS-over-RDMA), POSIX standards for parallel I/O, PVFS2, iSCSI, iSER, SRP, and DAPL; HPSS version 6 deployment with full Linux

support and no DCE dependencies, version 7 development with small file improvements, and continuing development support for the transfer agent, PFTP, and LDAP integration; and a security guideline for pNFS.

It is expected that most of the FY07 R&D of parallel file system and I/O libraries will yield candidate technologies for future petaFLOPS systems. These technologies, coupled with the most promising advanced data transport technologies and the petascale network environment candidate technologies, will be assessed for inclusion in a petascale system.

#### ***1.5.4.5.7 Advanced Networking (SNL)***

The Advanced Networking project provides next-generation networking technologies in the 10–100 gigabit per second performance range spanning the range from system interconnect networks to the wide area environment. Advanced capabilities, such as RDMA, are investigated to enhance ASC computational throughput and efficiency. Security, problem diagnostic capabilities, and performance monitoring of network services are also investigated and developed as part of this project. The project coordinates R&D foci and effort with our tri-lab partners, academia, and industry to influence the development of next-generation technologies that meet the most demanding ASC requirements.

Petascale computing platforms require connectivity to other petascale resources such as storage, visualization, and archive systems. A network that scales up to the performance of these platforms is required to maintain a balanced, efficient system. This project will define the requirements, technologies, and architectures that will scale the network up to the petascale range. The network will be deployed to support the petascale systems before their implementation.

ASC large-scale Linux clusters are rapidly growing in number across the tri-labs to meet ASC capacity computing needs. Increasingly, these clusters use InfiniBand for their inter-process communication fabric, and in the near future are expected to use InfiniBand for the I/O traffic as well. The project provides national leadership in InfiniBand's evolution as an HPC technology by influencing the development of the OpenFabrics/OpenIB common Linux software stack effort, as well as contributing to InfiniBand-awareness in the Open MPI messaging library.

The project is also focused on enhancing distance computing and machine network user availability through the development of an automated, proactive, performance monitoring system that encompasses the ASC Red Storm and WAN environments.

Planned activities in FY07 include collaboration with the tri-lab community to develop a petascale network environment roadmap; partnership with industry, government, and standards bodies to pursue gigabit Ethernet standards beyond 10 gigabits per second; contributions to the further development of the OpenFabrics/OpenIB InfiniBand and RNIC software stack and Open MPI messaging libraries with an emphasis on scalability issues, such as congestion control and dynamic routing; investigation into alternative network transport protocols, including the SCTP protocol, with respect to providing high performance data transport, such as RDMA over IP, in the petascale environment; collaborations with the tri-lab community on InfiniBand adoption in large-scale clusters and data centers at the 10–30 gigabit per second performance level; the development of network simulation models (OPNET) to identify critical network design parameters for the petascale environment; collaboration with the Red Storm production networking activity to plan the network transition to a future petascale environment; and the prototype evaluation and first phase implementation of an automated monitoring and analysis system.

Major FY07 deliverables include the completion of a petascale networking environment roadmap that will outline suggested research areas to address expected technology deficiencies; detailed testing and analysis of some of the critical technologies identified in the petascale networking environment roadmap in collaboration with our tri-lab partners; the further development of open source OpenFabrics/OpenIB InfiniBand and Open MPI software to support dynamic and quasi-dynamic routing, subnet management scalability, topology awareness, multicast enhancements for MPI collectives, and enhanced problem diagnostics and performance monitoring tools; and a proactive monitoring system capable of automated analysis of selected system parameters in the ASC Red Storm and WAN environment.

With InfiniBand aggressive technology roadmap beyond 100 gigabits per second in the next couple of years and the rapid maturation of other petascale network technologies that will be identified in the FY07 roadmap activity, in FY08 the project will be focused on preparing for these technologies for the next-generation ASC petascale environment that will have to handle a diverse load of I/O and interprocess communication traffic.

#### ***1.5.4.5.9 File Systems and I/O (LANL)***

The File Systems and I/O Project provides end-to-end, high-performance networking and scalable I/O infrastructure for the ASC program. It also delivers high bandwidth, low latency interconnect technologies for the ASC compute platforms. The ASC program requires system and storage area network bandwidths at over 100 gigabytes per second, global file system I/O rates beyond 100 gigabytes per second, and latencies in the 1 microsecond range. All this performance must be provided in an integrated, usable, and secure way. Data transfer and storage bottlenecks are now a critical concern for current-generation, high-performance computing environments. Successfully meeting the ASC programmatic milestones requires carefully balanced environments in which the I/O infrastructure scales proportionally with increased ASC platform capabilities and application data needs.

This project is a coordination point for planning of all online storage, network, and data movement activities within the ASC program at LANL. These capabilities include online file systems such as the NFS complex and local supercomputer file systems, GPFS development, deployment and management, scalable I/O middleware development and support, interconnect technology development and deployment, and storage area networking development and deployment.

In FY07, this project will handle NFS support including planning for first generation scalable NFS and NFSv4 deployment, Panasas file system support, maintenance, regression testing, MPI-IO support, and enhancement for small and unaligned I/O, including development, testing, and deployment. Additionally, the project will extend the Panasas file system via the PaScalBB storage area network to provide for the new LANL Roadrunner capacity system. It will deploy OpenIB on new InfiniBand clusters, including the new LANL Roadrunner capacity system, implement the next phases of the PaScalBB to provide file system and archive bandwidth, and investigate alternative interconnects such as Quadrics, Myrinet, Infiniband, and 10 Gige. In addition, it will provide program relevant file system/storage metrics and support and encourage best I/O practices through publications of best practices presentation and support of code teams.

In FY08, the project plans a study of possible next generation archival technologies and direct attached file transfer agents. It also plans to work on Panasas scalable metadata deployment and a demonstration of RDMA data movement capability, and deployment of a full collaborative caching layer in MPI-IO. It also plans to demonstrate the next

generation InfiniBand technology and deploy NFSv4. The project will also investigate and deploy scratch file system capability for the petaFLOPS portion of the new LANL Roadrunner system.

### **WBS 1.5.4.6: Pre- and Post-Processing Environments**

This level 4 area provides integrated environments to support end-user simulation set up, and post-processing visualization, data analysis and data management. The scope of this product includes planning, research, development, integration and deployment, continuing customer/product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include tools for optimized problem set-up and meshing, metadata and scientific data management, and application-specific and general-purpose visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data. Special focus will be placed on tools for improving end-user productivity. Also included are procurement, deployment, and support of office and collaborative space visualization displays, mechanisms for image data delivery, and custom graphics rendering hardware.

#### **Pre- and Post-Processing Deliverables for FY07**

- Contribute to the deployment of a tri-lab ASC Purple environment.
- Prototype tools and/or techniques with applicability to V&V.
- Support the General Availability of Red Storm environment infrastructure.
- Apply the Design-through-Analysis tools to uncertainty quantification, beyond the traditional finite element analysis and to ensemble analyses.
- Custom visualizations of ASC users' simulations, engineering data, and other data as needed.
- Manage contracts for the visualization facilities and visualization software.
- Develop, as needed, for the planned unified tri-lab petascale data analysis tool: software development plan, quality assurance plan, configuration management plan, software requirements specification, and architectural design documents.
- Provide briefings in the facilities as needed.
- Train users for software and facilities use.
- Develop visualization software extensions that support parallel, progressive, and lossy compression techniques.
- Develop visualization software for analyzing and debugging ASC simulations by comparison and quantification of differences between ensembles of data using hardware accelerators.

#### **Strategy to Deliver Pre- and Post-Processing for FY07–FY08**

- Existing solutions for scientific visualization and data analysis are not common across laboratories. A key activity for FY07 will formalize a process to converge on a unified tri-lab VTK-based, open-source approach for a petascale data analysis tool combining the best capabilities of the VisIt and VTK/ParaView projects. This new tool will address critical petascale computing environment data analysis and data management needs, and is envisioned to co-exist with CEI's EnSight proprietary visualization tool.

- Develop and implement problem setup, data management, data analysis, and visualization tools to provide QMU-enabled comparisons of ASC simulation results with validation measurements.
- Coordinate the efforts of separate projects in areas such as scientific visualization, scientific data management, and systems and services integration to deliver the appropriate capabilities needed for new platforms such as Red Storm, Purple, BlueGene/L, and Roadrunner.
- Continue management and oversight of contracts with 1) Computational Engineering International, for customized development and support of the proprietary, production EnSight visualization tool suite; and 2) Kitware, for customized development and support of the open source ParaView/VTK system, which provides a foundation for active collaboration, prototyping and possible rapid delivery of advanced capabilities for the unified tri-lab petascale data analysis tool.
- Develop and deliver an integrated, efficient Design-Through-Analysis capability that spans the spectrum from pre- to post-processing for computational engineering.
- Leverage past investments in ASC VIEWS through further integration of visualization into the production environment.
- Examine uses of extensions of the present system, which hold promise to increase developer or designer productivity.
- Participate in planning efforts targeted at the eventual delivery of a petaFLOPS computing environment.

#### ***1.5.4.6.1 Scientific Visualization (LLNL)***

The Scientific Visualization Project is tasked with providing high-end rendering, compositing, and analysis capabilities for all levels of scientific simulation within the Defense and Nuclear Technologies (DNT) Directorate at LLNL. This can be broken down into three activities: 1) computer graphics research, 2) long-term exploration into data management and rendering techniques, and 3) custom visualization application development, providing tools that complement the scientists' existing toolset to provide expanded capability.

The team members of this project have extensive experience handling data processing needs of DNT users. Our researchers are world-renowned for their expertise with data and image compression, including on-the-fly compression of very large datasets; work on multi-resolution streaming architectures for data processing and visualization; topological "summary" techniques for scientific datasets; and view-dependent rendering and data management techniques. We have extensive experience with parallel rendering techniques, both for tiled displays and compositing architectures. The team exploits the latest capabilities of clustering hardware, graphics processing unit advances, and parallel storage systems.

There are several projects started in FY06 that will be extended in FY07. These include deploying our simulated radiography work as a library to be used by two DNT applications, continued development of the Chromium Render Server for remote desktop and tiled wall delivery, refinement of the high-dimensional analysis capabilities in VisIt, deploying and enhancing run-time visualization and analysis in VisIt, and extending our topology infrastructure to support new data analysis capabilities. We will continue to work toward delivering capabilities within VisIt to leverage its large installed user base. New activities that we will embark on in FY07 include adding floating-point compression to several IO libraries.

### ***1.5.4.6.2 Scientific Data Management (LLNL)***

The Scientific Data Management Project is focused on providing users with better ways to generate, organize, search, access, extract, compare, track, and archive large-scale scientific data. This is achieved through two major themes: production applications for enhancing data access and organization, and data discovery techniques and tools for representing, exploring and extracting pertinent information. The Scientific Data Management tools effort at LLNL develops production applications that provide solutions for everyday data management needs. Its goal is to present the user with an easy-to-use, consistent interface to data, regardless of location, and to automate and simplify repetitive tasks. Data discovery develops scalable algorithms for the interactive exploration of large, complex, multi-dimensional scientific data. It applies and extends ideas from data mining and pattern recognition in order to improve the way in which scientists extract useful information from data.

The Scientific Data Management tools team has expertise in data transport protocols, graphical user interfaces, web technologies, data representation, and advanced system architectures. This team developed the Hopper file management tool and the SimTracker simulation data management tool. Hopper features a sophisticated software infrastructure allowing seamless interaction with many file archival and transfer protocols, allowing the user to enjoy the same user interface no matter where the data resides or where Hopper is running. Hopper provides advanced data operations including synchronization, access checking, and searching by attribute and content. SimTracker provides a web-based view of a user's simulations, with powerful capabilities for archiving, documenting, and sharing the data within a workgroup. The data discovery effort focuses on the Sapphire data mining project, which has expertise in data mining, image analysis, video tracking, statistical techniques, and pattern recognition. The algorithms (some of these have been issued patents) and software in Sapphire are used for the comparison of simulations to experiments, the analysis of experimental images to extract statistics for use in building models, and feature extraction and tracking in data from high fidelity simulations of fluid mix problems, utilizing Sapphire capabilities in pattern recognition, advanced image processing, noise removal, feature extraction, and feature tracking.

FY07 activities in the tools area will focus on addressing key user requirements, including: providing a way to visually compare files on different hosts, providing optimized utilities for managing data in parallel file systems, enabling users to easily annotate and find simulation results, and extending the client/server model in our file management system to support scripting and inter-language use. Sapphire will continue its activities in data analysis, addressing problems of interest to DNT scientists. Specifically, it will continue the analysis of experimental images and simulation data, the analysis of high-fidelity fluid-mix simulation data for use in building models, and the analysis of data from molecular dynamics simulations. It will also continue enhancing the algorithms available in Sapphire and deploy the software as needed.

For Scientific Data Management tools in FY08, we plan to provide tools and methods for understanding and visualizing information gleaned from huge numbers of files in archives and elsewhere. The goal is to help users better understand and navigate massive numbers of files. In Sapphire, we plan to continue analysis of experimental and simulation data as requested by DNT scientists.

### ***1.5.4.6.3 Systems and Services Integration (LLNL)***

The Systems and Services Integration Project is focused on three main areas: visualization servers, displays, and operational support of the pre- and post-processing

product resources for ASC computing customers. The server effort includes architectural evaluations and planning, testbed prototypes, component and system testing, and computer security planning and execution, in addition to the deployment, system administration, system debugging, and maintenance of visualization servers. The display effort includes high resolution and high performance display devices and tiled display arrays for user offices and collaborative use areas. Also included are the necessary switching systems and control devices. Operational support includes managing the use of five Powerwall facilities and the associated servers, running video production labs, and applying and consulting on software including resource management tools, movie players, animation and visualization packages. The project also includes liaison activities with customers, ensuring that tools and capabilities are meeting user requirements and that they are easy to use.

The project has expertise in systems administration, software design and development, a wide range of visualization hardware and software tools and techniques, advanced displays and related infrastructure, windowing systems, graphics processing units, video production, computer animation and 3D modeling, and user interfaces. Hardware capabilities include five production visualization servers with a variety of central processing units, interconnects, graphics cards, and disks. The Linux clusters are connected to a Lustre file system and video display infrastructure that drives Powerwalls and smaller displays. We install, maintain, and consult on a wide variety of software tools (including EnSight, VisIt, IDL, Tecplot, AVS, and NCAR). We consult with users on our visualization infrastructure tools including movie players (Blockbuster, xmovie), visualization resource management software (telepath), software to run tiled displays as one logical display (DMX), and other tools (ViSUS, TeraScale Browser, and Hopper). The project supports demonstrations on the PowerWall, and we have the capability to visualize simulation data and create movies to be shown on Powerwalls or laptops. The project maintains both an unclassified and classified video production lab, each of which includes desktop systems with video editing software, 3D modeling and animation tools, and an assortment of video peripherals. We create DVDs and videotapes as needed in addition to supporting live demos.

FY07 activities include integration of a new visualization server in our testbed as part of our ongoing architectural evaluations of cluster technologies. We will continue to maintain all of the visualization servers and advanced displays currently supported, and to provide operational support for all visualization facilities. Priorities include applying visualization tools and techniques to the large ASC data sets that will be generated on BlueGene/L, Purple, and Peloton. We will continue to provide consulting and support for graphics and video production, and we will keep the production graphics software environment current and consistent. Funding permitting, we will procure and deploy a delayed B-111 projector upgrade. We will also continue our evaluations of image delivery technologies and display products.

Planned FY08 activities will include conducting ongoing architectural evaluations and planning for future visualization systems and displays. We will procure and integrate a visualization platform to support the visualization of data from new platforms. We will continue to provide consulting and support services in graphics and video production, and we will continue to keep our production graphics software environment current and consistent on the high-performance platforms, including the compute servers.

#### ***1.5.4.6.4 Capacity Visualization (LANL)***

Visualization and visual analysis are essential tools needed by code teams and designers in understanding the terabytes of data that are generated in a single simulation run. The Production Visualization Project, sometimes referred to as "Production Visualization,"

provides visualization services from the machine to the desktop for users in the ASC program. Most importantly, people funded by Production Visualization work with code teams and designers to visualize their datasets, train them so that they can visualize their own datasets, assist them in using the large facilities, and assist in giving briefings in the facilities to high-level visitors. Software currently supported at full strength by the Production Visualization Project includes EnSight, EnLiten, EnVideo, Vista, PoP and GMV. This project also supports and maintains the large facilities, including the new CAVE, the PowerWall Theater and the co-laboratories. It maintains the visualization infrastructure, which delivers video from the machines to the users' desktops. Finally, it assists in the process of bringing new machines up by troubleshooting graphics systems on these machines, by performing the visualization software integration tasks needed, and by installing and maintaining critical visualization software on ASC machines.

This project is the place where all client-oriented visualization work takes place in the ASC program at LANL. Capabilities include:

- Custom visualizations for Code Teams A and B, ESA Division and RRW
- Expert knowledge of visualization software, such as EnSight, EnLiten, Vista
- Development and support of PoP and GMV
- Support of the large facilities, such the CAVE and the PWT, including assistance, training and developing new tools to work with these facilities
- Briefing support
- Deployment and maintenance of Viz Corridor software, such as Matrix

In FY07, this project will assist in the transition of Viewmaster from deployment to full production status. It will continue to maintain the visualization corridor, both hardware and software. Finally, it will continue oversight of the facilities and software maintenance contracts.

In FY08, this project will continue to support ASC visualization facilities and respond to users' visualization needs.

#### ***1.5.4.6.5 Capability Visualization (LANL)***

For ASC simulation and experimental application teams visualizing, analyzing, and understanding their results is the key to providing answers to questions of national interest and advancing scientific understanding. These activities are hampered by the need to visualize ensembles of massive datasets possibly stored on remote supercomputing resources. The Capability Visualization and Analysis project works to overcome these barriers to effective visualization and analysis through the creation of the next-generation visualization software and hardware technology. Our hardware effort identifies, exploits, and integrates state-of-the-art technologies in graphics, networking, and computer architecture to create a visualization hardware system infrastructure that overcomes these barriers. Our software effort creates visualization software tools that use this hardware system infrastructure to address ASC user requirements. As the infrastructure and tools mature, they will be deployed to the users and incorporated into the production environment.

In FY07, this project will develop solutions to the remote/distance visualization problem. Traditional scientific visualization techniques display images only after receiving all of the data needed for the entire image. Progressive visualization techniques incrementally improve a displayed result over time. For example, web browsers progressively load an extremely large image showing an increasingly detailed

result over time. We will build upon our VTK/ParaView visualization software infrastructure so that it supports parallel, progressive, and lossy compression techniques. ASC scientists will be able to see results quickly and progressively without having to wait for long periods of time for entire results to be visualized or transferred over the network.

In FY07, this project will also develop a hardware system infrastructure based upon graphics accelerators (for example, GPUs and IBM Cell processors) and a corresponding new software infrastructure that supports the comparison and quantification of differences between ensembles of data using hardware accelerators. The software will incorporate an analysis language for expressing derived fields and comparison metrics. It will support automatic visual comparisons and integrated quantitative measures of properties of the data. This software will be useful both for analyzing and debugging ASC simulations.

In FY08, this project will continue its mission of creating next-generation visualization software and hardware technologies that meet ASC user needs. We will build upon these deliverables in the distance visualization area to automatically analyze datasets and prioritize the transfer of data over the network. Methods to support user definition of scientific features will be incorporated into the graphic accelerators/comparative visual analysis work.

#### ***1.5.4.6 Scientific Visualization (SNL)***

The Scientific Visualization project researches, develops, deploys, and supports software that, together with associated infrastructure, delivers a data analysis and visualization environment in support of ASC platforms and applications. This project develops capabilities to meet the most difficult post-processing challenges facing the ASC community, whether at a desktop computer or in a high-resolution collaborative environment

A major thrust for this project has been the development and support of scalable visualization software within the open source ParaView/VTK framework that makes effective use of the aggregate power of cluster-based platforms. In addition, the project employs ParaView to deliver Sandia-developed advanced capabilities in support of Sandia application requirements such as parallel volume rendering for highly unstructured mesh data and the correct visualization of high-order unstructured elements. This same open source framework will now also provide a common basis for the emerging tri-lab petascale data analysis tool, which will have a stronger emphasis on data analysis capabilities. Toward that end, the project is developing core capabilities in ParaView 3.0 that enable scalable quantitative and comparative analysis and advanced support for scripting. Together with customizable interfaces, these core strengths will support targeted interfaces and capabilities for key SNL user communities, including V&V and High Energy Density Physics (HEDP).

Computational Engineering International's EnSight software is a proprietary commercial application, which has been enhanced in partnership with this project. This project supports EnSight's use at Sandia, as a baseline production scientific visualization tool with which several hundred Sandia users are familiar.

The project develops and/or provides certain other software and/or utilities as needed. Tools are also designed and implemented to enable access to high performance resources from the desktop. The project provides testing, benchmarking, installation, and end-user support for its full suite of tools and is working with other elements of the pre- and post-processing environment product to deliver an integrated DTA capability.

Planned activities in FY07 include support the deployment of the tri-lab ASC Purple user environment; provide targeted tools for V&V and HEDP visualization needs provide direct Sandia user support, especially as needed for the ASC Purple user environment; and participate in planning for the unified tri-lab effort to develop advanced data analysis capabilities as a continuation of the ParaView / VTK effort, including those capabilities targeted at the eventual delivery of a petaFLOPS computing environment.

Preliminary planned activities in FY08 include development and support of advanced data analysis and visualization tools within the tri-lab toolkit; continued development of customized/targeted interfaces; continued participation in the development of an improved DTA environment; and continued direct user support.

#### ***1.5.4.6.7 Scientific Data Management (SNL)***

The Scientific Data Management project equips users with effective tools with which to access, move, search, extract, compare, track and generally post-process large-scale scientific data. The metadata and data tools effort focuses on enhancing data access and organization, while the data discovery effort develops techniques and tools to represent, explore, and extract pertinent information.

At Sandia, the metadata and data tools effort provides easy to use components for processing and analyzing simulation data. These components ease large data movement, management, and manipulation. Examples are dsacp (large data movement between HPC systems and HPSS archive), zephyr copy (also known as zcp, large data movement between HPC systems), and dsabrowse (browse and manage the HPSS namespace). Metadata and Data Tools also provides tools for data manipulation. The primary tool is dstk, a Python based interactive program that permits data derivation, subsetting, extraction, output, and plotting.

The data discovery effort provides a collection of techniques and tools for the detection, representation and extraction of information from simulation data. Tools in the data discovery suite include the ParaView attribute editor, algorithms and software for rapid, robust "commodity" pattern recognition in terascale data, ParaView Lookmarks (lookmarks are to data browsing what bookmarks are to web browsing), the FCLib feature characterization software library (for building characterizations of regions of interest that eliminate the tedious parts of post-processing analysis), the variety of example custom characterizations built with FCLib, and software to support rich detection of the significant differences between data sets.

In FY07, in the metadata and data tools context, one focus will be generalized code verification in support of HEDP and other codes. This effort will include mechanisms to input and generate reference solutions, computation of mesh metrics, error metrics, and convergence rates, and support for various forms of output for computed data. Another focus will be advanced data translation capabilities, particularly enhancing mesh variable mapping to support coupled electrical/mechanical simulations. Metadata and data tools will also continue to support its current capabilities and its integration with DART and visualization tools.

In FY07, one data discovery focus will be the use of the mature FCLib library to support and build tools for the comparison and contrast of characterizations across data sets, even across data with distinct meshes, and to apply these tools in support of V&V sensitivity analyses. On the terascale pattern recognition front, we will down-select from current candidates and implement in the AVATAR Tools software a single technique for handling partitioned training data, which is statistically dissimilar. The end result will be an end-to-end set of tools that permit a user to mark example regions of interest in a

data set, then automatically find and characterize all such regions in all of the simulation variations run as part of a sensitivity analysis, automatically compare each run with the others, and efficiently visit, via Lookmarks, points of interest detected by the comparisons.

In FY08, the data discovery effort will extend the FCLib based-comparison capabilities to support cross-code, as well as cross-mesh, data comparison. It will also re-factor appropriate parts of the AVATAR Tools suite to provide massive data pattern recognition capabilities on the desktop, not just on the large ASC machines.

#### ***1.5.4.6.8 Infrastructure Deployment and Support (SNL)***

The Infrastructure Deployment and Support project develops, deploys, and operates data analysis systems and display facilities. The project provides terascale systems that will interactively analyze, visualize, store, and archive output from the largest scientific computers: Sandia's capacity-computing clusters, Red Storm, ASC Purple, and beyond. This project deploys capabilities to meet the most difficult post-processing challenges facing the ASC community, while also ensuring the availability of computing and storage systems that enable effective analysis of simulation results by ASC environment customers.

A major effort has been the deployment of the Red RoSE cluster in support of classified simulation on Red Storm, the deployment of an upgraded Feynman cluster to provide initial support of unclassified simulation on Red Storm, and demonstration of the institutional file system capabilities. Highlights of the past year include Limited Availability of the Red RoSE and Feynman clusters in support of Red Storm, General Availability of Red RoSE, record breaking performance of RoSE's parallel file system, and record breaking performance of RoSE's parallel visualization capability.

Principal activities planned for FY07 are acquisition, deployment and general availability of a black RoSE cluster; deployment of the institutional file system; and support and maintenance of Red RoSE. The black RoSE cluster will replace the aging Feynman cluster and will support unclassified simulations on Red Storm, Thunderbird, and Sandia's other capacity computing systems. The black RoSE cluster is essential to General Availability of the unclassified Red Storm Environment. The institutional file systems will provide high-bandwidth shared access by Sandia's large computer systems to RoSE's petascale storage systems. This will enable the output from large simulation runs to be written directly to a common storage system, shared by the batch computing systems and the interactive pre- and post-processing computer systems, avoiding the need to transfer data from one system to another. The institutional file system is designed to be a cost-effective way to meet both the large storage capacity needs of batch computing and the high-performance capabilities required for highly interactive visualization and data analysis.

FY07 activities will also include basic operations and maintenance of existing capabilities in Sandia's ASC-supported collaborative visualization facilities.

Preliminary planned activities in FY08 include replacement of the older parts of Red RoSE and an upgrade of its capabilities to support the expanded Red Storm Environment and the Purple Environment, as well as ongoing support of existing infrastructure.

#### ***1.5.4.6.9 Analysis Management and Prototyping (SNL)***

The Analysis Management and Prototyping project provides a finite element analysis environment that enables analysts to significantly reduce the time required to complete

the DTA process. The DTA environment comprises a federation of analysis setup and post-processing tools that interoperate through a common metadata standard. These tools help analysts rapidly complete tasks such as CAD model geometry editing, mesh building and editing, application of mesh parameters (for example, boundary conditions and materials information), input deck creation, job submission, post-processing and visualization of analysis results, and archiving of analysis artifacts. To reduce analysis time, DTA technologies focus on three general areas of opportunity: reducing the time required to traverse individual process steps, reducing the fraction of rework required during subsequent visits to process steps, and reducing the likelihood of backward process iteration.

In FY07, the DTA Realization Team will upgrade the initial limited capability environment releases and will implement support for new capability needs identified by the weapons development community. General environment improvements include a more comprehensive and integrated means of archiving the many artifacts produced during the course of an analysis, with the goal of being able to reconstruct a specific analysis as needed and being able to reuse artifacts from previous analyses. DART will also implement an automated regression testing mechanism for the environment and will institute a nightly test procedure to rapidly identify technical problems as individual tools release upgrades. The team will also provide broader support for ASC compute platforms and applications.

The Sandia weapons community completed a modeling and simulations requirements analysis and identified needs of the problem setup environment. Two primary themes from that analysis are an increased emphasis on uncertainty quantification to support engineering design decisions and on greater use of analysis tools beyond the traditional mechanical and thermal finite element tools. Specifically, electrical system, radiation transport, and electro-magnetic modeling were identified as high priorities.

DTA will provide comprehensive support for ensemble analyses, which is needed to quantify analysis uncertainties. Specifically, DART will implement a data passing mechanism with DAKOTA to support the use of DDACE for running predetermined sets of simulations as well as to support design optimization. DTA will also begin to support important non-finite element analyses, such as electrical system, radiation transport, and electro-magnetic modeling.

#### ***1.5.4.6.10 Model Generation Tools (SNL)***

The Model Generation Tools project provides analysts with advanced capabilities for creating and manipulating complex simulation models. The project develops tools that edit model geometry to prepare it for finite element meshing; generate, augment, and edit meshes; provide access to pedigreed materials information; and generate input decks for the Sierra codes. Specific tools include CUBIT, SIMBA, and Materials/WISDM. CUBIT provides a wide variety sub-tools that support geometric manipulation to prepare a model for meshing and for creating the meshes themselves. SIMBA provides sub-tools for modifying existing meshes, augmenting them with additional information, managing collections of model artifacts, and producing input decks for many ASC applications. Materials/WISDM is a materials information system and data manager.

During FY07, the CUBIT project will continue to enhance its virtual geometry capability, adding automation sub-tools that will facilitate cleanup and preparation of solid models for mesh generation. The enhanced virtual geometry capability will significantly reduce the amount of work an analyst must perform to generate a meshable geometry. The CUBIT team will also develop infrastructure to support multiple geometric models for design optimization studies and uncertainty quantification. This capability will include

support for geometry and mesh parametric changes as well as additional support for geometry-based boundary conditions.

Materials/WISDM will expand parameter set estimation and evaluation capabilities and incorporate the Lamé Library of material models. By doing so, it will expand the number of supported material models by an order of magnitude. Materials WISDM will also enhance its user interface, building on the results of extensive interviews with the analyst community performed during FY06. Finally, Materials WISDM will enhance interoperability with other DART tools to eliminate manual data translation and resulting data errors.

SIMBA development efforts will focus on three areas. The input deck management capability will be expanded to provide comprehensive support of input deck features for all of the Sierra codes. Based on feedback from the analyst community, the mesh editing and manipulation capabilities will be modified to present a visual select/operate paradigm. Finally, SIMBA will provide new tools to examine both a mesh and an input deck to detect problems in the simulation before the job is submitted, thereby ensuring that the correct simulation is run.

## ***WBS 1.5.5: Facility Operations and User Support***

### **WBS 1.5.5.1: Facilities, Operations, and Communications**

This level 4 product provides necessary physical facility and operational support for reliable production computing and storage environments. The scope of this product includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities and collaborations. This product also covers physical space, power and other utility infrastructure, and LAN/WAN networking for local and remote access, as well as requisite system administration, cyber-security and operations services for ongoing support and addressing system problems.

#### **Facilities, Operations, and Communications Deliverables for FY07**

- Ensure power and cooling sufficient for expanded Red Storm system.
- Free up floor space of approximately 2000 square feet in the Central Computing Facility.
- Reduce presence of “mixed mode” rooms in the Central Computing Facility (such as, consolidate systems of like security levels).
- Deploy NFSV4 server for production use (Q4FY07).
- Release Kerberos 1.4.X and OpenSSH for distribution (Q4FY07).
- Document process for moving DFS data to NFSV4 servers (Q3FY07).
- Deploy an additional pair of 1-GE encryptors installed by all tri-lab sites to allow 4 Gbps file transfers.
- Deliver 10-GE Type 1 IP encryptor production qualification plan.
- Deploy Sandia remote access to ASC Purple.

- Deliver a report delineating the different scenarios available for efficient usage of the ASC WAN bandwidth.
- Deploy an IPv6-capable ASC WAN.
- Complete the 10-G infrastructure to full 50 10-GE paths on both the classified and unclassified environments.
- Complete the system-wide communications performance improvements for the local environment to extract maximum performance from the 10-GE network infrastructure investment.
- Provide world-class support for the local 10-GE infrastructure, including automated testing and notification of performance issues.
- Complete the integration of the capacity and capability network systems to enhance seamless transition of data and projects between them.
- Ensure the availability of the WAN and HPC computing environment (for example, the network, platforms, applications, and services).
- Increase capacity for long-term hierarchical storage through phased replacement of existing silo equipment with newer technology.
- Increase memory per node on Red Storm to restore balance of original system. Anticipate goal of 6 gigabytes per node or 4 Gigabytes per core using Dual Core Opteron processors and 400 MHz DDR memory. Complete upgrade of all Red Storm nodes to Dual Core Opteron processors.
- Deliver system utilization, reliability, and availability reports to HQ demonstrating the best-in-class characteristics of Red Storm.
- Improve infrastructure to supply file system storage for HPC clusters and Viz/Data services platforms.
- Initiate acquisition of increased storage capacity for Sandia California's centralized file storage systems and deployment of a common user environment across all production compute systems. A prototype Lustre file storage system will be implemented that will be shared by NWCC and visualization servers.
- Upgrade the electrical capacity in the SCC to support Roadrunner (level 2 milestone due June 2007) and follow-on HPC systems planned for the SCC in out years. In addition, there will be several upgrades for the network and I/O infrastructure also in support of the Roadrunner level 2 milestone.

#### **Strategy to Deliver Facilities, Operations, and Communications for FY07–FY08**

- Coordinate activity with 1.5.4 projects engaged in obtaining resources and implementing changes to the computing infrastructure. Consolidate some system operations activity by reducing the variety of platforms supported. Identify procurement actions required in FY07–08 to support long-term data storage needs for scientific systems. Expand code support infrastructure to enhance software quality practices and test and build procedures for ASC codes. Continue test and evaluation activities with high-speed encryption device providers while implementing additional encryptors into tri-lab WAN.

##### ***1.5.5.1.1 System Administration and Operations (LLNL)***

The System Administration and Operations project provides for the ongoing system administration and computer operations functions for the successful management and

support of the ASC platforms and computing environment. It includes laboratory-maintained hardware support capabilities.

Capabilities include highly skilled system administration to ensure installation, integration, and ongoing support of ASC platforms including operating system and software configuration; feature, functionality, and security patches; troubleshooting, analysis, and diagnosis. This project also includes 24/7 operational monitoring capability for unclassified and classified computing environments consisting of large-scale computing platforms, infrastructure components, and networks. For laboratory-maintained systems, this project provides hardware maintenance capabilities including component inventory and replacement.

In FY07, we plan to complete integration and initiate ongoing support activities for Purple and continue ongoing support for BlueGene/L. We will complete the integration and initiate ongoing support of a new White replacement capacity platform. We will continue to train and develop staff in laboratory-maintained hardware support.

Preliminary planned activities for FY08 include planning and preparation for next generation petascale computing platforms.

#### ***1.5.5.1.2 Software and Hardware Maintenance, Licenses, and Contracts (LLNL)***

This project provides for vendor-provided hardware and software maintenance, support, licenses, and development contracts.

Capabilities include negotiated hardware and software maintenance and license contracts to ensure a robust ASC computing environment and to protect the computational investment of the NNSA. Targeted development contracts to enhance the capabilities of specific software components are also included.

Planned activities in FY07 include ongoing vendor support for designated platforms.

Preliminary planned activities for FY08 include ongoing vendor support for designated platforms.

#### ***1.5.5.1.3 Computing Environment Security and Infrastructure (LLNL)***

The Computing Environment Security and Infrastructure project provides for computing environment infrastructure services, software, servers, workstations, and desktop systems necessary for the efficient, effective, and secure operation and support of large-scale ASC platforms.

Capabilities include ongoing integration, development, and support on unclassified and classified networks of robust infrastructure environments to support large-scale ASC platforms including but not limited to name and time services, backups, staff productivity tools (for example, e-mail and messaging), cyber security tools and technologies.

Planned activities in FY07 include continued ongoing support of currently deployed services and capabilities. Investigation and analysis of new technologies for the possible replacement of existing technologies will also continue.

Preliminary planned activities for FY08 include continued ongoing support of currently deployed services and capabilities. Activities also include investigation and analysis of new technologies for the possible replacement of existing technologies.

#### ***1.5.5.1.4 Facilities Infrastructure and Power (LLNL)***

The Facilities Infrastructure and Power project provides for the necessary physical facilities, utilities, and power capabilities to support staff and the ASC computing environment.

Capabilities include adequate raised floor space, cooling facilities, and power to site large-scale ASC platforms. In addition, needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff is part of this project.

Planned activities in FY07 include facilities upgrades necessary for the siting of an initial delivery petascale platform in late 2007 or early 2008. Upgrades to the TSF include a 3 MW expansion to bring it up to the original design point of 15 MW for computing and the completion of the mechanical cooling system in the building. Also planned is the ongoing maintenance and support of existing computational and staff facilities, and continuing analysis of future ASC computing requirements to anticipate the needed modification and/or expansion of facilities to be able to site next generation ASC platforms.

Preliminary planned activities in FY08 include ongoing maintenance and support of existing computational and staff facilities, and analysis of future ASC computing requirements to anticipate the needed modification and/or expansion of facilities to be able to site next generation ASC platforms.

#### ***1.5.5.1.5 Classified and Unclassified Facility Networks (LLNL)***

The Classified and Unclassified Facility Networks project provides the architecture design, planning, procurement, deployment, and operational support of the classified and unclassified facility networks.

Capabilities include a thorough understanding of the resource deployment roadmap acquired by participating in ongoing facility-wide planning efforts that includes the archival storage, visualization, platforms, capacity computing, and file systems. Network design, procurements, and deployments are updated and scheduled to accommodate these plans and ensure the network connectivity, performance, reliability, security, and operational support is available for the facilities to meet the requirements of all subsystems is also part of this project.

In FY07, we will complete the integrated planning process so the networking requirements are understood, particularly to meet our requirements and accommodate major deployments for new system procurements and further deployments of the Lustre file system. These major deployments present a formidable challenge to provide significant 10 Gigabit Ethernet infrastructure within a severely constrained budget. Further, to contribute to the effort to establish common tri-lab HPC RAS definitions and establish a reference instantiation, each lab recently established tools in operations that will provide improved metrics on the I/O infrastructure and configuration management of the network.

Planning for the facility networks in FY07 and beyond is an ongoing process and will be similar to other years. The actual network activities are dependent on the requirements. A network challenge in FY08 will be optimizing the performance of and managing the 10 Gigabit Ethernet infrastructure.

### ***1.5.5.1.6 Wide-Area Classified Networks (LLNL)***

The Wide-Area Classified Networks project provides the architecture design, planning, procurement, deployment and operational support of the classified wide-area networks, namely the DisCom WAN and the SecureNet WAN.

Ongoing discussions with the tri-lab community are critical to this project to ensure the network requirements for those remote users and facilities are mutually agreed upon and understood. This project must also plan far in advance to ensure the required NSA Type 1 encryption products are available, since these products are not commercial and have a long R&D and product development lead time. Operational support of these wide-area networks also requires effective and regular communication with and cooperation between the tri-lab network support teams. These activities will help ensure the proper planning occurs for the WANs, and the operational support is effective for the broader tri-lab user community.

Preliminary planned activities in FY07 strive for on-going improved reliability, responsiveness to problems, and application performance. With new platforms expected at LLNL and SNL, the DisCom WAN bandwidth may be inadequate to meet the tri-lab user requirements. This would require an upgrade to the network components and encryptors to provide this bandwidth.

### ***1.5.5.1.7 Requirements and Planning (LANL)***

This project includes planning activities for computing operations, collection and statistical evaluation of user requirements for computing resources, and development of new metrics and data collection.

The primary capability of this function is to collect and understand user requirements for production computing resources and quality of service, and to develop new metrics and data collection and analysis techniques to assist these purposes.

In FY07, this project will focus on building the infrastructure and support necessary to collect the raw data and distribute analyses in an effective and scalable manner, coordinated with the collection, analysis, and dissemination of system-level RAS and operations correctness data. FY07 activities will also include the planning and development of new metrics and/or measures for continuous improvement in quality of production services, especially with regard to coordination and correlation of user-level with system-level data.

FY07 will also see coordination of a system for definition collection and reporting of RAS data with the other NWC laboratories.

In FY08, we will continue to contribute to improvements in the quality and reliability of computing requirements collection and analysis, both at Los Alamos and in the tri-lab arena.

### ***1.5.5.1.8 Roadrunner Base System Integration and Operations (LANL)***

The focus in this area is to standup the delivered Roadrunner connected units, perform system tests on the overall compute system, and integrate the system into the Los Alamos computing environment for production work.

Completing the system integration is a major level 2 effort for Los Alamos. It defines the overall system environment delivered to end users for running application codes. Once completed, users can compile and run programs, use the system debugger, the parallel

file system, the archival storage, and the scheduling system to perform simulation and analysis of weapons systems that are suitable for these systems.

In FY07, we will integrate all hardware and software components to deliver a system environment to application users for programmatic work. Integrated quality assurance testing will be performed to ensure production readiness of systems. The core team to deploy Roadrunner will reside in FOUS. The team will focus on the following areas:

- System/OS
- File systems
- Interconnect
- External network including PaScalBB
- Regression test
- Monitoring
- Application readiness

The objective is for Los Alamos to deploy and test the Phase 1 Base Capacity system and get selected applications ready to use the system. The Base System will be delivered in FY06 and will be deployed in FY07. The Base System consists of 14 Connected Units (CUs) for classified computing, one CU for unclassified computing, and an Advanced Architecture Initial System (AAIS) for advanced system and algorithm work for unclassified computing. The AAIS includes a full CU with 56 cell processors. Each CU consists of 144 nodes. The 14 CUs for classified computing will be delivered to and installed in the SCC by the computer vendor. The other CU and the AAIS will be delivered and installed in the LDCC.

The 14 CUs in the SCC will be initially configured in the yellow unclassified network in preparation for the acceptance testing by the vendor. It is Los Alamos' intention to then move the system into the secure environment. Most of the CUs will be used for access by a limited set of applications for the rest of the fiscal year. A limited number of CUs will be used for integration testing in preparation for the installation of the cell blades in the Base System cabinets (if the Option for Phase 3 is exercised).

The short-term objective is for Los Alamos to deploy the Roadrunner system focused mainly on meeting requirements in capacity computing. Beyond the base system, however, this system is architected so as to form the basis for an advanced architecture computing platform.

Other strategies include using an open resource to test or validate early system modifications, using PANASAS for the global parallel file system for Linux (which has been in operation at Los Alamos for a number of years), using a change control process, and using 10 Gig-E for external IO.

A key strategy for transitioning the system to production in a smooth manner is for management and members of the production computing staff to be directly involved on the integration aspects of the system. The system integration team will form a peer relationship with the vendor staff to implement the technical and schedule tasks identified in the Statement of Work. The Production computing team will form a customer/vendor relationship with the vendor to provide an integrated support structure for the system for all operational issues. The areas currently identified for coordination include:

- Resource Management
- Near-line storage

- HPSS/PSI
- User Training
- User documentation
- Security Plans
- User environments/tools
- Account management
- Consulting Readiness
- Off-hours support

### ***1.5.5.1.9 Ongoing Network Operations (LANL)***

Ongoing Network Operations encompass all aspects of the network infrastructure that support ASC computing in the classified and, to some extent, the unclassified arena. This includes directly attached networks to HPC systems (machine area network), network backbones, the users' local area network, and the high-end wide area network (DisCom WAN) connecting the tri-labs.

The network aspect of production computing provides the technical expertise and operation required to support high-end production computing and data storage services for the use of ASC computing. FY07 commitments include continued network support for ASC Roadrunner deployment, Q, Lightning/Gordon/Bolt, Panasas global file system, and HPSS and access, via the DisCom WAN, through which users at Los Alamos gain access to Red Storm at Sandia, and Purple and BlueGene/L at Livermore. Network services in the classified partition are required to promote and maintain an environment that is acceptable to all users in order to complete their assignments. Increased availability and security require more automatic network monitoring and intrusion detection in the classified network. Increased network backbone bandwidth within the Los Alamos campus is being implemented to support growth in HPSS storage from local and remote ASC computing and improve data delivery for scientific visualization. Increased bandwidth over the DisCom WAN is being implemented for anticipated requirements from Los Alamos users of the new remote systems and Sandia and Livermore. Thrust areas consist of:

- Network support: Designing, developing, acquiring, deploying, and supporting classified network hardware and services to support computational systems. This includes operating and maintaining services such as e-mail, authentication, and Web servers, plus operating and managing the high-performance network backbone, services networks, local paths to ESN, SecureNet, DisCom and the WAN, the high-performance parallel interface/gigabit and ten gigabit networks.
- WAN operation and support: In close cooperation with Livermore and Sandia, operate and maintain the 10 gigabits per second WAN, previously known as the DisCom WAN. Deploy new technologies to increase utilization, bandwidth, and reliability where warranted. Deploy new internet protocol encryptors to position WAN growth and control costs. Plan for bandwidth growth from two to four times current bandwidth if use by Los Alamos designers warrants such growth. In addition, new vendor offerings such as Geomax WAN services from Qwest have provided cost savings over 2004 levels of up to 40 percent for the WAN link from Los Alamos to Sandia.
- Work closely with the Integrated Cyber Security Initiative project to deploy the new Enterprise Secure Network (ESN), which will not only include SecureNet operations

but will also address directory, email, and other network services currently running locally at each site or over SecureNet. Plan for a demonstration of the new ESN working with an ASC application in the future when ESN is deployed.

Planned activities in FY07 include continued operation and maintenance of the LAN/MAN/WAN infrastructure that in the past was covered by Production Computing Network Support and the DisCom WAN project. The LAN/MAN planned activities are to complete deployment of 10GigE technology in high traffic areas and provide network expertise in the deployment of the PaScalBB (Parallel, Scalable, and BackBone for a global file system) effort. In the WAN area new internet protocol encryptors will be deployed to position the WAN for future growth if and when necessary. Cost of the tri-lab WAN link has dropped as long distance data links are becoming more affordable.

For FY08, the activities will include continued operation and maintenance of the LAN/MAN/WAN infrastructure. We also anticipate more integration with the ESN deployment. We expect further deployment and tuning of the I/O infrastructure to support Roadrunner as it increases in capability.

#### ***1.5.5.1.10 (Network) Infrastructure Integration (LANL)***

Infrastructure Integration encompasses the entire network infrastructure upgrades needed to support integration of new HPC computing platforms planned for FY07 and beyond. This includes designing, procuring, prototyping, testing, and installing network switches and routers needed to support the new HPC platforms and provide increasing bandwidth to HPSS, VIZ systems, and customer workstations. In addition, this will also include integration of the network components into the HPC computing platforms and storage systems. This project will also include integrating new switches/routers into our network monitoring platforms and training operational personnel to monitor and troubleshoot the new systems.

In FY07, Los Alamos will integrate the new Roadrunner capacity system into the Integrated Computing Network infrastructure. We will extend our GPFS networks, PaScalBB, to support new HPC clusters and I/O links in both the Yellow and Red networks. We will also upgrade network infrastructure, where necessary, to increase bandwidth to HPSS, VIZ platforms, and customer workstations. We will install new backbone routers that support single and multiple 10 GigE trunks between HPC platforms and customer LANs. We will also provide level three development and support for system interconnect reliability and performance management.

Activities for FY08 include involvement in developing and evaluating a petascale environment for the computing environment with emphasis on the network technology that will enable this environment envisioned at LANL in FY08 and beyond.

#### ***1.5.5.1.11 Ongoing Systems Operations (LANL)***

This project includes all services for computational systems operated by Los Alamos for the purpose of developing and running user codes. These services include system configuration, computer security, resource management, system administration, system operation and monitoring. The project works with users to troubleshoot problems experienced with running their applications on new systems, and also works with users to plan and carry out transitions off of older platforms. Thrust areas include:

- Computer system support (on-going): Conduct daily system administration and continual monitoring of production systems and infrastructure servers. Install operating system patches and other software. Monitor systems for security events.

Analyze and improve performance of key subsystems, including file systems. Develop scripts and tools to aid in administration and operation. Design and deploy new infrastructure architectures and services. Maintain security plans and tests for existing systems and create plans as required for new production resources.

- Scheduling environment daily management (on-going): Provide daily administration, monitoring, and problem resolution of software subsystems used in resource scheduling, including accounting and reliability data generation and reporting.
- RAS (on-going): Continuously improve the end-to-end level of service as seen by the users. Conduct ongoing studies and improvement projects in the stability of large, integrated systems, including the development of improved diagnostic and monitoring capabilities.
- RAS metric collection, reporting, analysis, and management (FY07 effort): There are a number of RAS-data related projects underway or needing to be started in HPC division. Components:
  - Redesign RAS database to transition away from Remedy
  - Participate in tri-lab common RAS state diagram design and implementation
  - Participate in tri-lab common RAS analysis design and implementation;
  - Assess and select methods for obtaining RAS data for Bproc systems and for Roadrunner
  - Assess and select methods for obtaining RAS and throughput data for HPC networks and interconnects
  - Design and implement coordinated database for RAS, workload characterization, and tool regression testing data
  - Implement selected RAS monitoring tools for Bproc and Roadrunner systems
  - Assess and select and/or design data distribution and GUI for operations staff for BPROC and Roadrunner systems
  - Design and implement GUI for operations staff interaction with newly implemented RAS database
- System Decommission: Phase out and decommission systems at life cycle end. This includes terminating service contracts, sanitizing systems, and properly and securely disposing of media and hardware according to policy.
- Q System Decommission (FY07 effort): This transition overlaps critical milestones for the W88 MAR and the W76 LEP, which continue to be dependent upon the Q machine to date; care must be taken not to interrupt the completion of this work. Within these constraints, a phased transition of the Q machine out of service will be undertaken.
- System Operations: Ongoing operations currently requires staff to provide around-the-clock operations and monitoring of the scientific computing resources including high-performance computers, such as ASC Q and Lightning, and data storage and retrieval systems such as the HPSS. System operators monitor all components 24/7 and take corrective action and/or escalate problems to others who can effect corrective actions.

- **System Self-Maintenance:** Replace failing field replaceable units in selected systems on a 24/7 service basis. Repair field replaceable units and test to maintain a stock of reliable spare field replaceable units. Return failed components to suppliers as required. Maintain adequate stock of spare parts based on past and projected future component failure rates.

This project utilizes the skill, experience, and talents of professional teams of system administrators, operators, and other skilled technical resources. These resources utilize automated tools and documented procedures to administer, maintain, schedule, and operate all HPC systems and components. This includes all components of the production computing environment, including compute engines, networks, file servers, archival storage systems, the facilities they reside in and utilities they are dependent upon, and all required software on these systems.

Planned activities in FY07 include:

- Prepare to support the Roadrunner capacity system in production
- Plan for and execute the retirement of systems that have reached end of life
- Continue to provide system support
- Operations and management of all ASC production assets
- Study the feasibility of reduced staff (“lights out”) computer operations
- Implement additional formal and informal customer satisfaction metrics and measurement techniques
- Expand self maintenance to additional systems and components
- Enhance RAS metric collection, reporting, analysis and management
- Deploy a limited Tripod capability
- Q machine phase-out

Preliminary planned activities in FY08 include:

- Support the base Roadrunner system and advanced architecture components in production
- Prepare additional capacity systems for production and support them in production environment
- Continue to provide system support, operations and management for current systems
- Continue to provide system support, operations and management
- Implement a reduced staff (“lights out”) computer operations model
- Implement improved RAS reporting
- Deploy a more extensive use of Tripod tools and capability

#### ***1.5.5.1.12 Systems and Network Maintenance (LANL)***

This project obtains hardware and software support services for all production ASC systems and network assets operated by Los Alamos. Support services include vendor hardware and software support contracts; contracts for parts sparing and logistics for self maintained equipment; software licenses; annual software maintenance and support

fees; and contracts for telephone support and other remote hardware and software support requirements.

Two main activities are the initial procurement of services for new production resources, and the renewal of existing maintenance and support contracts, normally on an annual basis. This project also produces and maintains budgetary planning data and earned value reporting on hardware and software maintenance commitments and projects.

This project is a coordination point for all maintenance contracts for system and network hardware and software. Customers include ongoing operations and integration teams, projects and groups. Interfaces include third party vendors and suppliers and the Los Alamos supply chain and accounting organizations. This project, in coordination with supply chain, negotiates new and renewal contract pricing and terms and conditions. It initiated the completed in a timely way so that lapses in support and disruptions in services do not occur. Orders, both original and renewal, are tracked as they proceed through procurement and fulfillment.

Planned activities in FY07 include refining order entry and tracking capabilities including interfaces with customers, supply chain, program management, and accounting; forging even stronger relationships and understandings with customers and external interfaces; developing methods and procedures to ensure the capture of maintenance requirements for all new production resources well in advance of their required date; procuring and ensuring fulfillment of all maintenance services required by existing and new systems in FY07; and planning for all maintenance services required in FY08 and beyond.

#### ***1.5.5.1.13 Systems and Network Equipment (LANL)***

This project ensures the procurement of systems and networking equipment required for support of ongoing production systems. In addition, procurements of related software support and third party labor are coordinated and executed through this project. This project also produces and maintains budgetary planning data and earned value reporting on systems and networking equipment procurement activities.

This project is a coordination point for system and network equipment and related software procurements. Customers include ongoing operations and integration teams, projects, and groups. Interfaces include third party vendors and suppliers and the Los Alamos supply chain and accounting organizations. This project, in coordination with supply chain, negotiates purchase pricing and terms and conditions. Procurements are tracked as they proceed through procurement and fulfillment.

Planned activities in FY07 include refining order entry and tracking capabilities including interfaces with customers, supply chain, program management, and accounting; developing even stronger relationships and understanding with customers and external interfaces; developing methods and procedures to ensure the identification of requirements for all new production systems and network equipment well in advance; procuring and ensuring delivery of all systems and network equipment required by existing and new systems in FY07; and planning for all equipment required in FY08 and beyond.

For FY07 this project will be executed in two phases:

- Identify and plan FY07 operational system and network equipment requirements and develop a time phased spend plan, consistent with the budget allocation
- Execute time phased plan; procure and ensure delivery of all required operational system and network equipment

Preliminary planned activities in FY08 include procuring and ensuring delivery of equipment required in FY08 and planning for all equipment required in FY09 and beyond.

#### **1.5.5.1.14 Ongoing Facilities (LANL)**

This project is in direct support of the two major deliverables: *2.4 MW SCC Infrastructure Upgrade Project* and *Limited Availability of Roadrunner Capacity System for Selected Stockpile Applications*.

Los Alamos' infrastructure support is handled by three projects: the SCC, the LDCC, and the CCF.

The SCC, LDCC, and CCF projects require staff to operate and maintain the mission-critical electrical and mechanical systems that support high-performance computing systems located in the three computing facilities. The project covers the cost of running the SCC facility including labor, materials and supplies, programmatic equipment and installation upgrades, electrical power, and all maintenance contracts associated with the critical equipment. It also covers the cost of running the programmatic portion (such as, labor, materials and supplies, and maintenance contracts) of the other two facilities, the LDCC and the CCF. The project provides support for infrastructure design upgrades, project management, user interface and oversight, and computer site preparation for new platforms for all three facilities. This work is in direct support of the tri-lab compute resource needs for meeting their ASC milestones for the SSP. Because the tri-lab community requires the systems to be operational at all times, the staff provides on-call support after hours and on weekends to ensure continued operations of the electrical and mechanical systems. The computer facilities were established to house large, mission-critical computer systems in support of the SSP and other research efforts. This project provides the ability to ensure the continued operation of current computing systems and support infrastructure, as required, to meet future programmatic computing needs. A part of this function is ensuring that facility site preparations have been made and design of electrical and cooling requirements have been met for planned computer systems. This project's main customer is the production systems team, which is responsible for providing compute resources necessary for SSP.

In FY07, planned activities include maintenance and support for ongoing operations including site preparations for future computer systems. In all of our projects, we will continue to work with the ASC Platform Project director, the Production Systems team, and the HPC user community to successfully meet the goals of the project. FY07 projects include:

- **Ongoing Operations.** This project ensures the continued operation of current computing systems and the support infrastructure as required to meet programmatic computing needs. The goal is to minimize interruption to the operation of systems and optimize the performance, reliability, and availability of the high-performance computers. It includes ensuring preventive and corrective maintenance for all electrical and mechanical systems, minimizing power outages, and identifying, designing, coordinating, and overseeing infrastructure upgrades to keep the facility operational as well as support future generations of supercomputing. On call support is provided on a 24/7 schedule to ensure continued operations of electrical and mechanical systems.
- **Site preparation to support new systems.** The project includes site preparation including design and installation of electrical and mechanical equipment, cable trays, tiles, and as-built drawings, for second-phase Roadrunner computer upgrades as well as for future computer systems planned by the division. This project is in

support of future platform acquisitions that will bring over a thousand trillion operations per second (a petaFLOPS) of computing resources into the tri-lab computing centers. Additional facility upgrades will be needed in the LDCC (for example, cooling) to support future growth in the unclassified computing environment.

In FY08, planned activities include continued operations and maintenance of electrical and mechanical systems for ongoing programmatic computing including site preparations for the third phase of additional Roadrunner computer installations.

#### ***1.5.5.1.15 Power, Security, and Facilities Support (SNL)***

This project includes power, physical security, and facilities support as well as cyber security operations. The first task provides power, access control, maintenance, and construction management services for efficient operation of central computing facilities. It also coordinates all projects impacting delivery of new equipment, removal of excess computing systems, modifications to facilities, and scheduling of actions affecting the cooling or power environment of said facilities. Cyber security operations provides the cyber security support for the ASC Program, including deployment and management of cyber security components related to authentication, authorization, and remote access. This task also has responsibility for maintaining security policies and procedures for SecureNet.

The power, physical security, and facilities support tasks includes management of vault type rooms, consultation on issues dealing with power and cooling, administration of access and escorts for tours and work details, point of contact for construction work or building renovations, and funding of electrical utilities expense. The cyber security operations task is focused on providing the authentication and authorization services used by applications for the purposes of remote access and data movement across ASC sites. This task has responsibility for maintenance of the software used as by the ASC authentication and authorization infrastructure as well as remote access.

#### **Power, Physical Security and Facilities Support**

Planned activities in FY07:

- Begin preliminary coordination activity to plan for building expansion in FY08
- Pay all operational charges for electrical service, physical security compliance, personnel safety training and refresher courses (electrical safety, chilled water processes, delivery of heavy or outsized equipment)
- Implement improved space planning and utilization of computing floor areas through a realignment of existing systems into more common security realms, providing a greater degree of flexibility and agility to respond to new space requests
- Address power and network limitations currently present in the CCF

Expected deliverables in FY07:

- Ensure power and cooling is sufficient for expanded Red Storm system
- Free up floor space of approximately 2000 square feet in CCF
- Reduce presence of “mixed mode” rooms in CCF (such as, consolidate systems of like security levels)

Preliminary planned activities in FY08:

- Initiate building expansion project for Super Computing Annex aimed at doubling floor space to 40,000+ square feet
- Support technology refresh of cluster computing systems, providing deactivation and removal planning and coordination in addition to staging and installation of new platforms

### **Cyber Security Operations**

Planned activities in FY07:

- Continue software and hardware support for the Kerberos, Directory, and DFS infrastructure:
  - Software maintenance of current releases of Kerberos, SSH, and support services such as the Group Information Service, Keytab Service, Directory Content Management Service
  - Kerberos upgrade with support for smart cards
  - Replace F-Secure SSH with OpenSSH based distribution; this will require changes to the current distribution mechanism and coordination with the other tri-lab security projects
  - Develop transition tools to move DFS data to NFSV4 servers.
- Provide the security officer for SecureNet
  - Maintain accreditation plans and policies for SecureNet
- Deploy NFS V4 server for production use in the ASC environment

Expected deliverables in FY07:

- Deploy NFSV4 server for production use (Q4)
- Release Kerberos 1.4.X and OpenSSH for distribution (Q4)
- Document process for moving DFS data to NFSV4 servers (Q3)

Preliminary planned activities in FY08:

- Shutdown all DCE and DFS services at Sandia (Q2)

#### ***1.5.5.1.16 Network Operations, Deployment, and Distance Computing (SNL)***

This project includes three tasks that together provide the network infrastructure to allow users to access ASC resources, within Sandia and remote to it, as if they were local. The project continues to enhance the distance computing and machine network infrastructure for sharing the tri-lab resources so that the customer will experience increased network availability and performance. The project finds innovative ways to continue to provide required services while enhancing efficiency. This project also provides the integration, testing, monitoring, and analysis of the ASC/DisCom HPC environment (for example, the network, platforms, applications, and services) and continues to develop, deploy, and standup the HPC networks that interconnect all of the classified and unclassified supercomputing resources within Sandia. This environment includes the capability and capacity supercomputers, storage, visualization, archival systems, high performance desktops, and the network interface to the tri-lab DISCOM environment.

This project provides first class support and operation of the Red Storm platform and tri-lab terascale infrastructure. It also supports the HPC network as part of a world-class deployment of over 500 10Gig Ethernet interfaces. This infrastructure enables the levels of performance required to interconnect some of the fastest computers in the world. The parallel data movement applications leverage the infrastructure to ensure the compute disks can be emptied faster than simulation data is written to them. The project will also enhance the distance computing and machine network user availability while improving monitoring and reporting of system resource availability. The project will continue development activities to keep pace with evolving customer requirements.

### **Tri-Lab Dedicated Wide Area Network Operations**

Planned activities in FY07:

- Maintain, operate, and manage the ASC WAN and ESNet links provided by the vendors (pay the bills)
- Increase the maximum customer available ASC WAN bandwidth, for a single data movement from 2 Gbps to 4 Gbps
- Obtain 10 Gbps IP encryptors engineering samples as soon as they are available for testing; investigate the opportunity to replace the parallel 1 Gbps encryptor installation the tri-lab WAN with single 10 Gbps encryptor units
- Enable Sandia customers to utilize the newest compute resource available by providing Sandia remote access to LLNL Purple
- Develop scenarios for aggregating bandwidth for efficient usage of the available 10 Gbps bandwidth
- Identify and upgrade the necessary ASC WAN elements to implement IPv6 to satisfy the U.S. government directive to support IPv6 on network backbones

Expected deliverables in FY07:

- Additional pair of 1 GE encryptors installed by all tri-lab sites to allow 4 Gbps file transfers
- 10 GE Type 1 IP encryptor production qualification plan
- Sandia remote access to LLNL Purple
- A report delineating the different scenarios available for efficient usage of the ASC WAN bandwidth
- IPv6 capable ASC WAN

Preliminary planned activities in FY08:

- Maintain, operate, and manage the ASC WAN and ESNet links provided by the vendors (pay the bills)
- Purchase and deploy 10 Gbps IP encryptors in the ASC WAN and decommission the parallel 1 Gbps encryptors; this will reduce the network complexity, increase availability, and performance
- Activate the IPv6 protocol in the ASC WAN satisfying the government mandate

### **High Performance Computing Network Deployment**

Planned activities FY07:

- Build out the rest of the 10 G network infrastructure

- Continue the system wide tuning of the parallel data movement applications
- Support end-customer transition to 10 GE in their networks
- Develop and initiate the plan to transition support of the HPC network to the production group
- Ensure FY07 purchases are aligned with the petascale network environment project
- Test new technologies that will improve the utilization of the 10 GE infrastructure

Expected deliverables FY07:

- Complete the 10 G infrastructure to full 50 10 GE paths on both the classified and unclassified environment
- Complete the system wide communications performance improvements for the local environment to extract maximum performance from the 10 GE network infrastructure investment
- Provide world-class support for the local 10 GE infrastructure, including automated testing and notification of performance issues
- Complete the integration of the capacity and capability network systems to enhance seamless transition of data and projects between them

Preliminary planned activities for FY08:

- Continue the plan to transition as much of the infrastructure support to the production group as possible while maintaining world-class support
- Support the deployment and tuning of the 10 G DISCOM bandwidth to ensure effective utilization
- Implement new technologies to improve utilization of the 10 G infrastructure
- Initiate deployment of the petascale network environment

### **Apps at a Distance: Integration, Operations, Status, Monitoring, and Testing**

Planned activities in FY07:

- Continued implementation and standardization of the tools to monitor, manage, analyze, and report on the ASC high performance WAN
- Enable customer utilization of Purple—the newest computational tri-lab resource—by ensuring availability of the ASC high performance WAN
- Increase HPC platform, network, application, and service status, monitoring, and analysis
- Purple integration into tri-lab status and monitoring
- Re-evaluation of user perspective testing: monitoring, status, and reports for the HPC environment

The expected major deliverable in FY07 is to ensure the availability of the WAN and HPC Computing Environment (for example, the network, platforms, applications, and services).

Preliminary planned activities in FY08 include continued integration of HPC environment elements into CEM/SOC monitoring, reporting, business service views, and user perspective testing.

### **1.5.5.1.17 Sandia Production Systems (SNL)**

This project coordinates and manages all the activities within the computing centers that support daily operations and scheduling of ASC capability and capacity computing systems, pre- and post-processing platforms and data management, and long term hierarchical data storage systems.

Sandia Production Systems in Albuquerque hosts the tri-lab expedited priority run meetings for Red Storm; negotiates system dedicated time, operating system upgrade schedules, and facilitate problem identification and resolution across the tri-lab user community; and manages account granting process user requests appropriate to the ASC priorities and system status of Limited or General Availability. The end-to-end computational process requires a variety of pre- and post-processing capabilities and a variety of communications, storage, archive, and infrastructure support systems. This project provides funding to manage the operation of the various commercial systems supporting ASC users—from the deployment, testing, and integration of systems into the user environment, to the continuing maintenance and management of systems in production operation. Most importantly, the project provides for the team of skilled and experienced technical staff that administers the complex HPC environment and supports system users.

California Production Systems provides funding for 1) maintenance and support of the California NWCC; 2) maintenance of local resources to visualize data produced on NWCC, ThunderBird, and Red Storm; and 3) support of the local data storage (NFS, HPSS, storage silos) available to analysis users in California. Our goal is to provide operation of stable, powerful, efficient, accessible, and easily useable computation and visualization systems for the California user community in both SRN (secure restricted network) and SCN (secure classified network) environments and for both R&D and production use.

#### **Sandia Production Systems, Albuquerque**

In FY07, we plan to continue to support General Availability computing needs of Red Storm tri-lab customers; track and resolve errors of omission or commission with respect to the computing environment and platforms; conduct monthly user seminars to disseminate “best practices”; advertise system schedules and provide a face-to-face feedback mechanism in addition to on-line contact methods; collect, analyze and distribute system utilization metrics and modify operational parameters to improve customer satisfaction; transition all customers to the use of the HERT usage forecasting tool for assigning run time on Red Storm and capacity clusters; and complete conversion of HPSS systems to Version 6.2 (delayed from FY06).

Deliverables in FY07:

- Increased capacity for long term hierarchical storage through phased replacement of existing silo equipment with newer technology.
- Increased memory per node on Red Storm to restore balance of original system. We anticipate a goal of 8 gigabytes per node or 4 gigabytes per core using Dual Core Opteron processors and 400 MHz DDR memory. (Complete conversion of all Red Storm nodes to Dual Core Opteron processors.)
- System utilization, reliability and availability reports to HQ demonstrating the best-in-class characteristics of Red Storm.
- Infrastructure improvements to supply file system storage for HPC clusters and Viz/Data services platforms.

Preliminary Planned Activities for FY08:

- Complete transition of hierarchical storage systems from Powderhorn to StorageTek 8500 models and highest density tape drives
- Plan acquisition of next generation capacity cluster systems

### **California Production Systems**

Planned activities in FY07 include continued support of all California production computing, visualization, and storage infrastructure; improved visualization cluster usage in combination with deployment of a user environment common with the NWCC systems; and close integration of visualization, storage, and compute resources. Furthermore, metrics will be captured for analysis to characterize compute usage, data flow, and resource usage for applications running on these systems.

Expected deliverables in FY07 include acquisition of increased storage capacity for Sandia California's centralized file storage systems and deployment of a common user environment across all production compute systems. A prototype Lustre file storage system will be implemented, which will be shared by NWCC and visualization servers.

Preliminary planned activities in FY08 include the installation of production quality LUSTRE file storage systems to be used by California compute and visualization systems on the SRN and SCN.

#### ***1.5.5.1.18 Applications in Support of Manufacturing Production and Connectivity (Y-12)***

This project will support the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. The project will include determination of optimal methods for connecting to ASC computing resources and job submission, execution, and visualization. The project will also support the implementation of a compute cluster at Y-12 to provide the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources.

FY07 planned activities include the evaluation and deployment of file transfer methods to efficiently move input and output data between Y-12 and the labs; develop network infrastructure plan to support large file transfer; utilize ASC cluster resource for Manufacturing Computed Tomography solution; and participate in Nuclear Weapons Complex ASC activities.

Expected deliverables include file transfer method for large files and results from Manufacturing Computed Tomography solution.

Preliminary planned activities in FY08 include continuing utilization of ASC resources for production manufacturing problems; implementing a plan to integrate Y-12 compute / visualization resource with ASC environment; and participating in Nuclear Weapons Complex ASC activities.

### **WBS 1.5.5.2: User Support Services**

This level 4 product provides users with a suite of services enabling effective use of ASC tri-lab computing resources. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. Projects and technologies include computer center hotline and help-desk services, account management, web-based system documentation, system

status information tools, user training, trouble-ticketing systems, and application analyst support.

### **User Support Services Deliverables for FY07**

- Provide user assistance and consulting via help desk and in person consultation; preparation and presentation of appropriate training sessions and documentation for tri-lab customers; coordinate account management activities among tri-lab users and affiliates; problem identification, resolution, and tracking; maintenance of historical records of problems and their resolution; reports on metrics such as system utilization and reliability (reliability metrics include interrupt and repair data from the individual components to complete systems); planning and user support for commercial production software tools; and direct end-user and application development support for effectively utilizing systems.
- Contribute to the HPC consortium community particularly for sharing common information such as training.
- Provide Wiki and Jabber (Q1FY07).
- Provide CLIK support and usage metrics (Q1FY07).
- Provide Virtual Library for HPC Training (Q1FY07).
- Provide Red Storm RAS metrics by node and job (Q2FY07).
- Write Red Storm for Dummies sample documentation (Q3FY07).
- Develop system and code models for use in predicting code behavior on future architectures.
- Improve code scalability on Red Storm and Purple for highest priority ASC codes.
- Demonstrate higher parallel efficiency on Sandia applications.

### **Strategy to Deliver User Support Services for FY07–FY08**

- Maintain a well-trained and motivated workforce capable of providing user support services for Purple, BlueGene/L, and any new systems.
- Focus on customer requirements for access and effective use of all ASC systems, take responsibility for direct customer service to local and remote users, and take responsibility for usage statistics and administrative interface for external users such as ASC Alliances.
- Achieve continuous improvement in the quality of user support and user satisfaction with all services through the use of appropriate metrics and user feedback mechanisms.
- Utilize quality software engineering design that is well planned and managed. We will be expeditious, incorporating code reuse when possible. Continue accumulation of knowledge base entries and lessons learned from early experience with Red Storm and Purple. Present results of scaling and benchmarking efforts to code developers to encourage adoption of optimum scaling in communications, computation, and I/O operations on each capability platform. Use model of codes and platforms to assist design choices in advanced architecture efforts and selection of capacity systems (Sandia).
- Refine and improve monitoring and status reporting of production platforms and increase interaction and off-duty support for batch production runs (Sandia).

- Continue close working relationship between Sandia central computing staff and the applications developers and analysts from Sandia, Livermore and Los Alamos who are using the bulk of Red Storm cycles. Improve collection of knowledge base items describing the best practices for Red Storm. Closely manage bug reports and schedule of fixes provided by vendor partners. Leverage experiences with Cray XT3 product line through continued interaction with Office of Science labs ORNL and LBNL, and the Pittsburgh Supercomputer Center.
- Continue to look for ways to provide better service to ASC computer users, including tri-lab users. Los Alamos will transition to a single help desk phone number for all computer-related issues. This will take some intricate planning and coordination in order to work well to support Los Alamos and tri-lab users.

#### ***1.5.5.2.1 Hotlines and System Support (LLNL)***

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC tri-lab computing resources.

This project includes computer center hotline and help-desk services, account management, web-based system documentation, system status information tools, user training, trouble-ticketing systems, and application analyst support.

Planned activities in FY07 include supporting ASC tri-lab systems in the following areas: provide daily technical consulting and user support services; respond to customer calls to resolve questions about systems use or status; provide account management activities and system event scheduling including technical bulletin publishing; develop expertise in and documentation for the new tri-lab resource management system; administer and maintain Remedy for problem reporting and tracking; continue development of the LCAMS product to improve efficiency and accountability of account management activities and policy implementation; investigate the possibility of using COTS Identity Management software as a replacement for LCAMS; develop and implement account management tools and procedures for supporting the Capability Computing Campaigns model of operation for Purple; and develop and maintain web pages, documentation, and training as needed.

For FY08, we expect to continue to provide the same activities and services described for FY07. In particular, we will begin to develop web pages, documentation, and training for new systems.

#### ***1.5.5.2.2 Integrated Computing Network Consulting, Training and Documentation, and External Computing Support (LANL)***

This project is comprised of three teams:

- The Integrated Computing Network consulting office, responsible for direct customer service to local and remote users of LANL ASC resources
- The training and documentation team, responsible for the development and delivery of documentation and training materials for LANL ASC resources
- The external computing support team, responsible for usage statistics and the administrative interface for external HPC users such as tri-lab ASC and ASC Alliances

Thrust areas consist of:

- User support services: Daily technical consulting and user support services. The Integrated Computing Network consulting office provides direct customer service to

local and remote users of the LANL HPC resources in support of ongoing computing and the ASC program. They respond to customer calls to resolve questions about systems use or status. They also conduct account management activities and system event scheduling. Additionally, this team provides administration and maintenance of user problem reporting and tracking systems.

- Web page development and support: Develop and maintain web pages at LANL in support of ASC.
- Operational metrics for HPC environment: Collect data regarding machine usage, machine availability, and mean-time-to-failure as needed by various interested parties and make that data accessible.
- Development and delivery of user documentation and training.

This project is staffed by teams of skilled and experienced consultants, trainers, documenters, and statistical professionals. This staff, coupled with a set of procedures and tools, delivers user support services including help desk, account management, documentation, and training. The staff also maintains statistics on system utilization and reliability.

Planned activities in FY07 include performing ongoing user support for users of Los Alamos' ASC computing resources; expanding class offerings, particularly in the use and support of Linux capacity systems and Panasas-based file systems; expanding online documentation in both breadth and depth; deploying new trouble tracking systems in the open and secure; deploying enhanced HPC statistical reports and infrastructure; re-architecting the accounts system; and furthering refinement and development of ongoing customer satisfaction surveys to quickly identify and deal with any future shortcomings.

Preliminary planned activities in FY08 include performing ongoing user support for users of Los Alamos' ASC computing resources and continuing to improve the quality of support for productive use of the computing services.

### ***1.5.5.2.3 Ongoing User Services (LANL)***

This project has three main focus areas grouped under the umbrella of user services. The first focus is to support the production environments by ensuring that tools needed by users are available and working on production platforms. Tools in this context include compilers, debugging tools, and performance and profiling tools. To do this, we research, develop, and deploy reliable, available, and scalable capabilities in commercial and custom software development. The essential elements must be identified and advanced to minimize cost, enable similar environments across like platforms, and enable developers to successfully develop, debug, and analyze applications. Specific activities include interacting with vendors and other third party tool providers to negotiate requirements, licenses and product issues; working with the project lead for Systems and Network Maintenance to budget for tool contracts and to forecast new requirements; testing tools on new computing resources before they go into production use; establishing acceptance/regression test suite for tools; assisting users in applying tools to their problems; continuing development of a comprehensive tools strategy, based on Los Alamos needs and in concert with the tri-lab tool workshops; and participating in planning and assessing tools environments for future ASC systems.

In FY07, support will be focused on the software stack that is vetted with tri-lab and local application teams. The supported products list has been established in FY06 to work toward better service for a smaller set of products. Another strategy is to keep the software versions consistent across the platforms.

Specific FY07 goals include:

- Various support levels will be provided depending on the product and needed support. At the low end, products will be installed and tested to ensure they are operational on platforms. At the higher end, development work may be included or coordinated to make the products work in our environment and to enhance for better utilization. User support and training will also be done at the higher levels.
- Continued interaction with users and vendors to resolve issues. With the scale and variation that we are running, this is a constant effort.
- Continued process improvement with HPC-7 teams for product installation, license server management, user issue, and support. The target is to get effective communication systems in place and partnerships for support.
- Software packages will be managed through a change control repository; RPMs will be made for installation where possible.
- A nightly (or periodic) regression test system for tools will continue to be implemented. This is to ensure the health of the software tools environment and track performance points.
- A monitoring and tracking system will check and compare tool versions running on platforms and segments. In addition, usage statistics will be gathered for some tools.
- Contract management for tool upgrades and support will continue.
- Third-party requests will be reviewed and supported once vetted with appropriate approval bodies. Some support will be provided to teams that want to install specific local tools.
- License server management will be increased with consistent architecture and backup license servers.

The second focus is on application development and end user in-person support. This involves technical consulting and collaborations with customers on issues related to the application code development environment and use on production systems, including documentation and training. The goal is to engage users directly to understand problems and issues and assist in their solution. The staff fulfilling this focus area will also serve as a two-way conduit for communication of requirements, issues, and information between users and all the ASC teams that support them.

Planned activities in FY07 include assisting users via direct support and with tools as the number and use of Linux capacity computing platforms grows and moves from 32-bit to 64-bit operation; providing support for the 64-bit production development environment on Lightning and new capacity clusters and establishing a focused approach to supporting users; ensuring an acceptance/regression test suite process is initiated for tool acceptance prior to production implementation; providing technical consulting and collaborations with customers on issues related to the application code development environment on production systems; and managing requests for and installations of third-party software on the production systems.

The third focus is on application readiness of systems. FOUS will coordinate and work with CSSE to deploy an applications readiness team at LANL that will work directly with users to address application problems with the ASC computers. This team is needed to support users/code teams to do initial ports and deep problem determination of code issues on new platforms. The intent of the team is to supply computer scientists that know both applications and systems. In FY07, a major focus of the team will be the

readiness of the Roadrunner Capacity System for the applications targeted for that machine.

Activities in FY08 include planning for software tools required for future systems and continuing to improve the quality of support for productive use of the computing services. This includes continuing the role of the application readiness team to help ensure issues with optimal use of ASC compute resources are being addressed.

#### ***1.5.5.2.4 Sandia Scientific Computing User Support (SNL)***

This project has two main focus areas: to create a change in the culture of HPC through knowledge management and to leverage code support activities for general computing needs of the ASC community. The project provides customer support to users of Red Storm by means of a web-enabled, online, 24-hour support process. The process includes access to a knowledge base and email support (CLIK), with additional telephone support available during regular weekday hours. It also includes porting assistance to new ASC architectures, tuning for scalability in processing and I/O, benchmarking performance of applications, performance modeling activity, and the system administration tasks of providing a robust, reliable, available computing platform. The project also provides resources for collaboration with other projects that support high performance computing users.

The web-enabled knowledge management form of support will facilitate collaboration among users and experts and will provide and archive the information automatically at the information owner's request. It will provide the most current information from distributed locations to allow users to access information that individuals need to do their job from one location.

Code-centered user support will supply expertise in trouble analysis, scaling efficiency improvements, application and middleware (library) interactions, compiler optimizations, file system and data transfer performance, benchmarking of applications and scaling studies to assist the rapid conversion of ASC applications to the newest capability platform (Red Storm and Purple in FY06 / FY07). This activity will benefit not only the applications developers but also the production computing system administrators and the high performance and wide area networking projects by supplying true application level workloads to the computer and the environment. The initial set of applications will be small and well understood, thus providing an enhanced ability to apply specific stresses to the systems for the purpose of identifying weaknesses or gaps in basic service and performance. Computational experts will consult on performance issues, trouble analysis and system configuration options, which, when combined, will result in enhanced throughput and efficiency of use for the capability platforms.

#### **Sandia Scientific Computing User Support**

Planned activities in FY07:

- Implement more collaborative areas, for example, Wiki and Jabber. Develop virtual library of HPC training available across sites with XT3 capability. Offer more classes via Grid Access and face-to-face classes.

Expected deliverables in FY07:

- Provide Wiki and Jabber (Q1FY07)
- Provide CLIK support and usage metrics (Q1FY07)
- Provide Virtual Library for HPC Training (Q1FY07)

- Provide Red Storm RAS metrics by node and job (Q2FY07)
- Red Storm for Dummies sample documentation (Q3FY07)

Preliminary planned activities in FY08:

- Plan transition of CLIK support application tool over to corporate supported programs

### **User Support – Code Centered Activities**

Planned activity in FY07:

Continue performance modeling and predictive performance work with additional ASC codes beyond the initial Limited Availability set of benchmarked and scaling study applications. Begin porting assistance to Purple and analyze performance and scaling characteristics of Purple platform with respect to Sandia engineering applications. Provide guidance to Sandia code developers and analysts using Purple and Red Storm on environment configuration choices that produce optimum runtime efficiencies. Assist all customers in analyzing error states and condensing error generating code into the minimal representation of the fault.

Expected deliverables in FY07:

- System and code models for use in predicting code behavior on future architectures
- Improved code scalability on Red Storm and Purple for highest priority ASC codes
- Demonstrated higher parallel efficiency on Sandia applications

Preliminary planned activity in FY08:

- Use system models to assist in evaluation of future advanced architecture options with respect to performance of Sandia applications
- Continued refinement of code performance on Red Storm Dual Core system and other ASC platforms (such as Purple)

### **WBS 1.5.5.3: Collaborations**

This level 4 product provides collaboration with external agencies on specific high-performance computing projects. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. This product also includes any programmatic support across the entire ASC program and studies, either by internal or external groups, that enable the program to improve its planning and execution of its mission.

#### **Collaborations Deliverables for FY07**

- Manage tri-lab contracts and procurements, support tri-lab meetings and conferences, and prepare and release ASC documents (Livermore).
- Provide communications and documentation support to headquarters (Sandia).
- Manage the Sandia/Russian science and technology projects (Sandia).
- Continue to support cooperative work with the SC06 ASC Research Exhibit and the FY07 Principle Investigators meeting (Sandia).
- Host of Predictive Science Panel (Los Alamos, Livermore).
- Consultant support to HQ for strategic programmatic planning and fostering of collaboration (Los Alamos).

### **Strategy to Deliver Collaborations for FY07–FY08**

- Support SC06 Research Booth.
- Provide staff and contractors to support Collaboration programs.
- Work with the other laboratories and HQ to meet deliverables on a timely basis.
- Work with other laboratories to organize and host the Predictive Science Panel meeting, to gather objective feedback and input concerning the predictive science capabilities of the National Program and to provide high-level guidance for planning (Los Alamos, Livermore).
- Provide contract personnel to assist with National Program planning and to assist with fostering collaborations (Los Alamos).

#### ***1.5.5.3.1 Program Support (LLNL)***

The Program Support project provides service to the ASC program. These services include document management, technical writing and editing, procurement and contracting, project management, and meeting support. These services are in support of both Livermore-only and tri-lab activities.

In FY07, this project will support preparation and release of technical reports, publicity media, and tri-lab documents such as FY07–08 Implementation Plan and publicity for the SuperComputing conference (SC06, to be held in November). Procurement activities will include management of existing tri-lab contracts and negotiating and executing new contracts. Meeting support will be provided for the FY07 principle investigator's meeting, SC06, Predictive Science Panel Meeting (expected to be held in April at Livermore), and other meetings and workshops.

In FY08, we plan to continue activities similar to those in FY07.

#### ***1.5.5.3.2 One Program/Three Labs (SNL)***

One Program/Three Labs funds several critical coordination and integration activities essential to the success of ASC. These are divided into four distinct parts: 1) development with HQ of a yearly plan for ASC communications and implementation of this plan; and 2) support for tri-lab and HQ interactions including ASC executive meetings, the ASC principal investigators meetings, and an SAIC contract for HQ administration support; and 3) Sandia outreach to the DoD laboratories; and 4) as part of the Russia/NNSA Science and Technology program, collaborative projects in modeling and simulations with the Russian nuclear laboratories and institutes. These crosscut and outreach activities seek to facilitate cooperation and collaboration amongst U.S. and international laboratories, improve program visibility within the high-performance computing community and enhance the overall operations of the ASC program.

Planned activities in FY07 include the continued support of the SNL and HQ ASC website, continued production of high quality communications materials for HQ and the broader high-performance computing community, support for the ASC Executive Committee, support for quarterly meetings of the ASC executive committee; support for the annual principal investigator meetings that expose attendees to technical and programmatic efforts at the three laboratories; management of the SAIC contract to provide various administration support to HQ, and semi-annual project review meeting with the Russian Federal Nuclear Laboratories and various National Academy Institutes.

FY07 activities will continue to be a tri-lab program management effort.

### ***1.5.5.3 Program Support (LANL)***

Los Alamos will continue to support collaboration within the ASC program by participating in the following:

- Alternate with Livermore in hosting the Predictive Science Panel. Results will be incorporated into Program plans and initiatives. The Predictive Science Panel consists of outstanding academic experts in the field of predictive science, and helps to build relations with the academic and industrial community of interest to predictive science.
- Provide consultant support to the federal program management efforts to foster collaborations and build support within the predictive science community.
- Act as lead in FY07 for organizing the annual tri-lab Principle Investigators (PI) Meeting.

## V. ASC Level 1 and 2 Milestones

Table V-1. Quick Look: Level 1 Milestone Dependencies

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
349	Deliver advanced ASC physics and engineering simulation capabilities to support the W76 and the W80 LEP/certification.	1	2006	Jun-06	ASC/Integrated Codes	HQ, LLNL, LANL, SNL	Enhanced predictive capability, DSW Deliverables, Increasingly productive environment	Yes	LANL-VV-06-01, LANL-IC-06-01, LANL-IC-06-02	Yes	359 (M44)
350	Provide a 100 teraflops platform environment to support to the tri-laboratory DSW and Campaign simulation requirements.	1	2007	Dec-06	ASC	HQ, LLNL	Enhanced predictive capability, DSW Deliverables, Increasingly productive environment	Yes	461, 464, 1348, LL-CSSE-06-02, LL-CSSE-06-04, LL-CSSE-06-05, CSSE-07-03	Yes	359 (and many more)
359	Complete modern baseline of all enduring stockpile systems with ASC codes.	1	2009	Sep-09	ASC/Integrated Codes	HQ, LLNL, LANL, SNL	DSW Deliverables		349		

**Table V-2. Quick Look: Level 2 Milestone Dependencies for FY07<sup>6</sup>**

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
2340	Performance assessment and improvements of hypre in ASC codes	2	FY07	Jun-07	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2341	Enhance capabilities of ASC simulation codes to support current and planned analysis efforts: ALE2D	2	FY07	Mar-07	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2342	Explore and assess new opportunities to enhance the predictive capability of ASC simulation codes: thermonuclear burn	2	FY07	Sep-07	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2343	Demonstrate a dynamic PEM model capability with multi-physics coupling in a performance code to support required DSW and RRW simulations	2	FY07	Jun-07	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2344	Provide improved 2D mesh generation capability to support DSW and RRW	2	FY07	Dec-06	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2345	Enhance capabilities of ASC simulation codes to support attribution efforts	2	FY07	Sep-07	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
2346	Develop a predictive HE detonation model for insensitive high explosives.	2	FY07	Mar-07	ASC/Physics and Engineering Models	LLNL	Enhanced predictive capability			Yes	MRT#359

<sup>6</sup> Factors such as FY07 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
2347	Next-generation Pu multiphase EOS delivered for QMU and stockpile certification studies.	2	FY07	Mar-07	ASC/Physics and Engineering Models	LLNL	Enhanced predictive capability	Yes	C206-01 C206-02	Yes	MRT#359
2348	Advanced hydro model implementation.	2	FY07	Jun-07	ASC/Physics and Engineering Models	LLNL	Enhanced predictive capability			Yes	MRT#359
2349	Develop and assess an advanced material strength model via a multiscale modeling approach and application of appropriate verification and validation tools.	2	FY07	Sep-07	ASC/Physics and Engineering Models	LLNL	Enhanced predictive capability			Yes	MRT#359
2350	K-Values Project	2	FY07	Sep-07	ASC/Verification and Validation	LLNL					
2351	Terascale Simulation Facility upgrade	2	FY07	Sep-07	ASC/Facility Operations and User Support	LLNL					
2352	Demonstrate capabilities in PREMO for current stockpile issues and eventual reentry in weather and hostile blast simulations	2	07	Sep-07	ASC/Integrated Codes	SNL					
2353	Demonstrate predictive science capabilities using Trilinos Release 8	2	07	Sep-07	ASC/Integrated Codes	SNL					
2354	Use of Aria to simulate laser weld pool dynamics for neutron generator production	2	07	Sep-07	ASC/Integrated Codes	SNL					
2355	Demonstrate capability to predict combined cable SGEMP, Box SGEMP, and cavity SGEMP events using the RAMSES modules CEPTRE, ITS, and EMPHASIS for RRW	2	07	Sep-07	ASC/Integrated Codes	SNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
2356	Develop high fidelity model of the SGT/AT accident analysis; and using the PRESTO code demonstrate simulation capability (including visualization) at scale on RedStorm/Purple capability computing platform	2	07	Sep-07	ASC/Integrated Codes	SNL					
2357	A new wide-range equation of state for tungsten with higher accuracy for warm dense matter	2	07	Sep-07	ASC/Physics and Engineering Models	SNL					
2358	Create and document a software library numerically implementing the response of advanced materials for solid mechanics that is linked to the ASC solid mechanics codes	2	07	Sep-07	ASC/Physics and Engineering Models	SNL					
2359	Demonstrate QMU methodology for hostile neutron-gamma environment in a stressing radiation encounter	2	07	Sep-07	ASC/Verification and Validation	SNL					
2360	Deliver prototype post-processing techniques and methodologies to support verification and validation (V&V) of SNL simulation applications.	2	07	Sep-07	ASC/Computational Systems and Software Environment	SNL					
2361	Capability upgrade of Red Storm	2	07	Sep-07	ASC/Computational Systems and Software Environment	SNL					
2363	Immersive Topology Modification Engine	2	07	Sep-07	ASC/Computational Systems and Software Environment	SNL					
2364	Red Storm (U) General Availability	2	07	Mar-07	ASC/Facility Operations and User Support	SNL	Increasingly Productive Environment				

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
2365	Release of a Crestone Project Code to Support RRW concepts, the W88 and TBI	2	07	Mar-07	ASC/Integrated Codes	LANL					
2366	Release of a Shavano Project Code to Support RRW concepts and the W88	2	07	Mar-07	ASC/Integrated Codes	LANL					
2368	Release of a Crestone Project Code to Support RRW concepts, the W88 and TBI	2	07	Sep-07	ASC/Integrated Codes	LANL					
2370	Release of a Shavano Project Code to Support RRW concepts and the W88	2	07	Sep-07	ASC/Integrated Codes	LANL					
2372	Implementation and validation of DSD for CHE and IHE	2	07	Mar-07	ASC/Physics and Engineering Models	LANL					
2373	ASC code development capability for attribution using radchem	2	07	Sep-07	ASC/Physics and Engineering Models	LANL					
2374	Implementation of new mix model for improved predictive capability	2	07	Sep-07	ASC/Physics and Engineering Models	LANL					
2375	Secondary V&V Assessment Supporting W76-1 LEP Certification (Final)	2	07	Dec-06	ASC/Verification and Validation	LANL					
2376	Define Primary Validation Test Suite Version 1	2	07	Sep-07	ASC/Verification and Validation	LANL					
2377	Engineering V&V Assessment of Truchas Code for an RRW Manufacturing Process	2	07	Sep-07	ASC/Verification and Validation	LANL					
2378	Limited availability of Roadrunner capacity system for selected stockpile applications	2	07	Jun-07	ASC/Facility Operations and User Support	LANL					
2379	Tri-Lab Study of Verification Methods	2	07	Sep-07	ASC/Verification and Validation	LLNL, SNL, LANL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
2380	Develop and deploy Tripod capabilities for capacity computing environment	2	07	Jun-07	ASC/Facility Operations and User Support ASC/Computational Systems and Software Environment	LLNL, SNL, LANL					

Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY08

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
LL-IC-08-01	Explore and assess opportunities to enhance the nuclear weapon simulation capability for performance code for future stockpile activities	2	FY08	Dec-07	ASC/Integrated Codes	LLNL	Increasingly predictive capability				
LL-IC-08-02	Evaluate opportunities for sharing of modules across ASC codes	2	FY08	Jun-08	ASC/Integrated Codes	LLNL	Increasingly predictive capability				
LL-IC-08-03	Develop enhanced capabilities in a multidimensional (1D,2D,3D) effects code to support current and planned effects modeling efforts	2	FY08	Mar-08	ASC/Integrated Codes	LLNL	Increasingly predictive capability				
LL-IC-08-04	Enhance capabilities of ASC simulation codes to support current and planned analysis efforts: thermal assaults	2	FY08	Mar-08	ASC/Integrated Codes (Joint with Campaigns 1 and 2)	LLNL	Increasingly predictive capability				
LL-IC-08-05	Explore and assess new opportunities to enhance the predictive capability of ASC simulation codes : ejecta	2	FY08	Mar-08	ASC/Integrated Codes	LLNL	Increasingly predictive capability				
LL-IC-08-06	Demonstrate ALE/AMR hydrodynamic capability in ASC codes	2	FY08	Jun-08	ASC/Integrated Codes	LLNL	Enhanced predictive capability				
LL-PEM-08-01	Assessment of nuclear physics uncertainties and requirements for improving the primary reactivity calculations.	2	FY08	Sep-08	ASC/Physics and Engineering Models	LLNL	Enhanced predictive capability			Yes	MRT#359
LL-VV-08-01	First Development of the Secondary Computational Assessment Project (SCAMP)	2	FY08	Sep-08	ASC/Verification and Validation	LLNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
LL-CSSE-08-01	Deploy Tri-Lab Petascale Data Analysis Tool and Infrastructural Support	2	FY08	Jun-08	ASC/ Computational Systems and Software Environment	LLNL	Increasingly predictive capability				
SNL-IC-08-01	Initial demonstration of failure capabilities in the SIERRA mechanics structural mechanics modules (PRESTO, ADAGIO)	2	FY08	Sep-08	ASC/Integrated Codes	SNL					
SNL-IC-08-02	Demonstration of heterogeneous combustion in the FUEGO module of SIERRA mechanics	2	FY08	Sep-08	ASC/Integrated Codes	SNL					
SNL-IC-08-03	Initial demonstration of neutron tube source plasma generation and transport	2	FY08	Sep-08	ASC/Integrated Codes	SNL					
SNL-IC-08-04	Next generation systems (advanced architecture follow on to #369)	2	FY08	Sep-08	ASC/Integrated Codes	SNL					
SNL-CSSE-08-01	Uncertainty Quantification Application Demonstration of Design-Through-Analysis Environment	2	FY08	Sep-08	ASC/ Computational Systems and Software Environment	SNL	Increasingly predictive capability				
LA-IC-08-01	Release of a Crestone Project Code to Support RRW concepts and the W88	2	08	Mar-08	ASC/Integrated Codes	LANL					
LA-IC-08-02	Release of a Shavano Project Code to Support RRW concepts and the W88	2	08	Mar-08	ASC/Integrated Codes	LANL					
LA-IC-08-03	Release of a Crestone Project Code to Support RRW concepts, the W88 and the TBI	2	08	Sep-08	ASC/Integrated Codes	LANL					
LA-IC-08-04	Release of a Shavano Project Code to Support RRW concepts and the W88	2	08	Sep-08	ASC/Integrated Codes	LANL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
LA-PEM-08-01	Improved MD and PIC simulation capability	2	08	Mar-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-02	3D Material model and Roadrunner simulations.	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-03	Implement new radchem diagnostic capabilities	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-04	Improved models for aging predictions.	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-05	Models & data for RRW	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-06	Opacity advances	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-PEM-08-07	Validation of new mix models implemented	2	08	Sep-08	ASC/Physics and Engineering Models	LANL					
LA-VV-08-01	V&V Assessment of ASC Codes for Thermonuclear Applications, including initial definition of secondary validation test suite	2	08	Dec-07	ASC/Verification and Validation	LANL					
LA-VV-08-02	Exchange of Primary Validation Test Suite with LLNL	2	08	Mar-08	ASC/Verification and Validation	LANL					
LA-VV-08-03	Validation Assessment of Penetration Mechanics for Surety Applications	2	08	Mar-08	ASC/Verification and Validation	LANL					
LA-VV-08-04	Catalog of Major Adjustable Parameters in Weapons Physics Simulations	2	08	Sep-08	ASC/Verification and Validation	LANL					
LA-VV-08-05	Engineering Validation Toolbox Tri-Lab Release	2	08	Sep-08	ASC/Verification and Validation	LANL					
LA-CSSE-08-01	Tri-lab Visualization Tool	2	08	Jun-08	ASC/Computational Systems and Software Environment	LANL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
LA-CSSE-08-02	Seamless and Productive Workflow	2	08	Jun-08	ASC/ Computational Systems and Software Environment	LANL					
LA-FOUS-08-01	Demonstrate Heterogeneous Computing	2	08	Mar-08	ASC/Facility Operations and User Support	LANL					
TRI-CSSE-08-01	Infrastructure Plan for Petascale Computing Environment	2	08	Jun-08	ASC/ Computational Systems and Software Environment	LANL, LLNL, SNL					

## Detailed Milestone Descriptions for FY07

<b>Milestone (ID#):</b> Terascale Simulation Facility upgrades (#2351)				
<b>Level:</b> 2				
<b>Fiscal Year:</b> FY07				
<b>DOE Area/Campaign:</b> ASC				
<b>Completion Date:</b> Sep-07				
<b>ASC nWBS Subprogram:</b> Facility Operations and User Support				
<b>Participating Sites:</b> LLNL				
<b>Description:</b> Complete a 3-MW expansion of electrical power for computing systems in the east computer room of the Terascale Simulation Facility (TSF) and complete the final connection of all air handlers that were delivered as part of the TSF line item.				
<b>Completion Criteria:</b> The completion criteria is to fully accept the system by testing, adjusting, balancing and starting-up the system in accordance with the detailed project commissioning plan to be developed by the system engineers during the design phase of the project.				
<b>Customer:</b> ASC Program				
<b>Milestone Certification Method:</b> Completion evidence will be documentation and Program hand-off.				
<b>Supporting Resources:</b> N/A				
<b>Codes/Simulation Tools Employed:</b> N/A				
<b>Contribution to the ASC Program:</b> Enables the TSF to site next generation ASC capability and capacity systems.				
<b>Contribution to Stockpile Stewardship:</b>				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Operations of the east computer room will be impacted during physical construction.	Medium	High	
2	The direct digital control (DDC) system commissioning of air handlers will require significant contingency planning to ensure existing operations are not impacted.	Medium	High	

<b>Milestone (ID#):</b> Deliver prototype post-processing techniques and methodologies to support verification and validation (V&V) of SNL simulation applications (#2360)				
<b>Level:</b> 2				
<b>Fiscal Year:</b> FY07				
<b>DOE Area/Campaign:</b> ASC				
<b>Completion Date:</b> Sep-07				
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment				
<b>Participating Sites:</b> SNL				
<b>Description:</b> The Pre- and Post-Processing product at SNL is in the process of developing data analysis tool features that will support V&V. In FY07, the product expects to deliver preliminary capabilities that enable comparative analysis of multiple data sets and the verification of simulation codes. Specific capabilities will be developed and refined to support specific targeted SNL V&V applications (under negotiation), which will include HEDP.				
<b>Completion Criteria:</b> Complete a set of V&V analyses that: (1) exercise the delivered data analysis capabilities on a range of application runs and/or data sets; and (2) are applied in direct support of SNL V&V application requirements.				
<b>Customer:</b> V&V program (Tim Trucano), HEDP applications				
<b>Milestone Certification Method:</b> A Final Report and/or other artifacts documenting results; Memo from customer accepting results				
<b>Supporting Resources:</b> Post-Processing development and support staff; V&V and HEDP partner resources to guide feature development and generate application data; computing resources on ASC platforms and/or data analysis servers as needed.				
<b>Codes/Simulation Tools Employed:</b> ALEGRA simulation code (HEDP); ParaView and possibly other data analysis/visualization tools.				
<b>Contribution to the ASC Program:</b> Provide important tools needed to enable V&V / QMU for ASC simulations in support of predictive capabilities.				
<b>Contribution to Stockpile Stewardship:</b> Provide important tools needed to support simulation-based weapon qualification/certification.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Inadequate access to V&V/HEDP partner resources.	High	Low	Medium
2	Inadequate access to simulation result data for V&V analysis; inadequate access to compute resources needed to do V&V simulation runs to produce results.	High	Low	Medium

<b>Milestone (ID#):</b> Capability upgrade of Red Storm (#2361)				
<b>Level:</b> 2				
<b>Fiscal Year:</b> FY07				
<b>DOE Area/Campaign:</b> ASC				
<b>Completion Date:</b> Sep-07				
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment				
<b>Participating Sites:</b> SNL				
<b>Description:</b> Complete the upgrade of Red Storm processors, and memory, to provide a larger capability platform and a consistent memory footprint for all computational nodes. Each node will be fitted with Dual Core Opteron processors and upgraded to 8 Gigabytes of memory. This upgrade will provide a consistent memory level for all processors within the compute partition of Red Storm, including the newly installed 5th row. Total compute node count will be 12,960 Dual Core AMD Opterons, with 103,680 Terabytes of memory. Service and I/O nodes will also be upgraded to Dual Core Opterons.				
<b>Completion Criteria:</b> Verification and Validation of installed hardware using base level hardware/firmware diagnostics followed by scaling studies reprising the earlier work using Sandia applications demonstrated during the General Availability Milestone effort.				
<b>Customer:</b> DSW, ASC V&V, ASC Integrated Aps				
<b>Milestone Certification Method:</b> Final report and internal review presentations of scaling study results.				
<b>Supporting Resources:</b> Sandia FOUS products and personnel				
<b>Codes/Simulation Tools Employed:</b> All tri-lab ASC and legacy codes				
<b>Contribution to the ASC Program:</b> Progress in meeting \$/Teraflop Program Goals				
<b>Contribution to Stockpile Stewardship:</b> Enhanced Capability for Modeling and Simulation				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
	Removal/replacement of modules reduces system stability	Medium	Low	medium
	Additional processor count creates anomalies in system software, reducing productivity/utilization	High	Low	medium

<b>Milestone (ID#):</b> Immersive Topology Modification Engine (#2363)									
<b>Level:</b> 2									
<b>Fiscal Year:</b> FY07									
<b>DOE Area/Campaign:</b> ASC									
<b>Completion Date:</b> Sep-07									
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment									
<b>Participating Sites:</b> SNL									
<p><b>Description:</b> Speeding up the time to generate an all-hexahedral mesh to facilitate simulation of complex nuclear weapons components is the focus of this milestone. The process of constructing an all-hexahedral mesh has been a significant bottle-neck in building models for simulation. To effectively address this issue, we propose to deliver a new Immersive Topology Engine for Meshing (ITEM) that promises to dramatically reduce the time to generate a hexahedral mesh.</p> <p>Currently the decomposition, feature reduction and topology editing necessary for a hexahedral mesh are expertise and time intensive operations. They require significant training, expert judgment, experience and control. Most users of existing hexahedral meshing technology use a trial and error approach to topology simplification and decomposition often resulting in work flow path dead-ends or reversals. The CUBIT geometry and meshing toolkit, developed at Sandia, is a technology leader in hexahedral meshing technology. As this tool has progressed we have seen a sharp increase in the number of novice and intermittent users. This has intensified the need for more automated topology modification mechanisms.</p> <p>The proposed Immersive Topology Engine for Meshing will be a new software tool built on CUBIT's strong CAD-based infrastructure. Given a complex weapon component or assembly solid model, the system will guide the user through the process of modifying the model topology so it is suitable for hexahedral meshing. The model preparation phase will involve the following immersive /interactive process:</p> <ol style="list-style-type: none"> <li>1. Automatically evaluate topology modification options</li> <li>2. Present a prioritized set of reasonable options with a preview of expected results</li> <li>3. Automate the execution of the chosen option with little or no user interaction</li> <li>4. Evaluate remaining topology repairs needed and repeat with a new set of prioritized modification options</li> </ol> <p>Once a topologically valid model has been defined, new many-to-many sweeping algorithms will be used to generate a hexahedral mesh. In addition, parallel hexahedral meshing technology will be used to generate high-resolution meshes of millions of elements. The following table illustrates the contrast between the existing, state-of-the-art in hexahedral meshing technology and what will be delivered in this milestone.</p> <table border="0"> <thead> <tr> <th><b>Current State-of-the-Art Hex Meshing</b></th> <th><b>Immersive Topology Modification</b></th> </tr> </thead> <tbody> <tr> <td>User driven.</td> <td>User controlled, but driven by geometric reasoning and algorithmic expertise.</td> </tr> <tr> <td>Expertise intensive.</td> <td>Novice friendly.</td> </tr> <tr> <td>Time intensive.</td> <td>Dramatically reduces user time.</td> </tr> </tbody> </table>		<b>Current State-of-the-Art Hex Meshing</b>	<b>Immersive Topology Modification</b>	User driven.	User controlled, but driven by geometric reasoning and algorithmic expertise.	Expertise intensive.	Novice friendly.	Time intensive.	Dramatically reduces user time.
<b>Current State-of-the-Art Hex Meshing</b>	<b>Immersive Topology Modification</b>								
User driven.	User controlled, but driven by geometric reasoning and algorithmic expertise.								
Expertise intensive.	Novice friendly.								
Time intensive.	Dramatically reduces user time.								

<p>Trial and error search paths with reversals common.</p>	<p>Smart impact prediction reducing the need for path reversals.</p>			
<p>No hints on remaining tasks or options. Next step is left to user expertise.</p>	<p>Progressively guided process, providing continuous help and feedback to completion</p>			
<p>Limited to one-to-one or many-to-one sweeping topologies.</p>	<p>Expanded coverage to include new many-to-many sweep topologies.</p>			
<p>Mesh resolution limited to capacity of single processor.</p>	<p>New parallel processing hex meshing algorithms dramatically reduce compute times and increase mesh resolution capacity.</p>			
<p>Deliverables:</p> <ol style="list-style-type: none"> <li>1. Immersive topology modification user interface built on CUBIT’s CAD-based infrastructure.</li> <li>2. Geometric reasoning algorithms to guide the user through the geometry preparation process.</li> <li>3. Geometric tools to decompose and modify the topology based upon the results of the geometric reasoning algorithms with user selection.</li> <li>4. Many-to-many sweep tool to generate a high quality hexahedral mesh on models where sweepable topology has been identified or successfully transformed.</li> <li>5. A parallel system for generating an all-hexahedral pave-and-sweep mesh.</li> </ol>				
<p><b>Completion Criteria:</b> Successful demonstration of automated topology modification for hex meshing on a selected set of weapons components. Demonstration of advanced sweeping and parallel hex meshing on these components showing improved throughput for generating hexahedral meshes on selected weapons component models. Relevant models will be selected in consultation with DART pilot projects.</p>				
<p><b>Customer:</b> SNL analysis line organizations represented by the DART Independent Review Team, NWC weapons and component designers.</p>				
<p><b>Milestone Certification Method:</b> A Final Review will be conducted and its results documented. Formal documentation, such as a report or a set of viewgraphs with a written summary, will be prepared as a record of milestone completion.</p>				
<p><b>Supporting Resources:</b> This effort is supported by the DART CUBIT Project task.</p>				
<p><b>Codes/Simulation Tools Employed:</b> CUBIT.</p>				
<p><b>Contribution to the ASC Program:</b> Has the potential for reducing the time to generate a model for simulation by 50 percent or more. This development will also positively impact our ASC throughput indicator scores for Sandia.</p>				
<p><b>Contribution to Stockpile Stewardship:</b> Significantly reduces the time for generating analysis models, making high end analysis far more plausible to address design concerns, and facilitating the use of modeling and simulation for fast turnaround SFI efforts.</p>				
<p>No.</p>	<p>Risk Description</p>	<p><b>Risk Assessment (low, medium, high)</b></p>		
		<p>Consequence</p>	<p>Likelihood</p>	<p>Exposure</p>

<b>Milestone (ID#):</b> Red Storm (U) General Availability (#2364)				
<b>Level:</b> 2				
<b>Fiscal Year:</b> FY07				
<b>DOE Area/Campaign:</b> ASC				
<b>Completion Date:</b> Mar-07				
<b>ASC nWBS Subprogram:</b> ASC / Facility Operations and User Support				
<b>Participating Sites:</b> SNL				
<p><b>Description:</b> Red Storm will deliver routine, reliable service to multiple users by providing fully functional system software and tools as defined in the Red Storm Usage Model Version 2.0. The capability provided through the previous Red Storm Initial Operation, Red Storm Limited Availability, and Red Storm Limited Availability User Environment milestones will be revalidated implicitly rather than explicitly in the process of delivering these general availability capabilities. Production operation and infrastructure to support the tri-lab user community will be demonstrated to the generally observed standard for ASC capability machines. Success criteria will include a set of measures showing a general trend towards increasing hardware and software reliability. An effective user support environment will be in place to allow operation and use of the machine to the generally observed standard for ASC capability machines. Completion criteria for the milestone will be derived from comparison with the Red Storm Usage Model Version 2.0 and corroborated by designers/analysts from each laboratory in the course of routine application execution. In addition, Sandia will submit a set of 3 performance scaling studies which will run under production conditions using ASC applications. A review of the milestone will be conducted in Q2FY07. Note that Version 2.0 of the Red Storm Usage Model includes a mapping of ASC tri-lab user requirements (as enumerated in the ASC Computing Environment Version 8.0) to capabilities being deployed in the Red Storm operational environment. Version 2.0 of the Red Storm Usage Model was released in Q2 FY06.</p>				
<b>Completion Criteria:</b>				
<b>Customer:</b>				
<b>Milestone Certification Method:</b>				
<b>Supporting Resources:</b>				
<b>Codes/Simulation Tools Employed:</b>				
<b>Contribution to the ASC Program:</b>				
<b>Contribution to Stockpile Stewardship:</b>				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure

<b>Milestone (ID#):</b> Limited availability of Roadrunner capacity system for selected stockpile applications (#2378)				
<b>Level:</b> 2				
<b>Fiscal Year:</b> FY07				
<b>DOE Area/Campaign:</b> ASC				
<b>Completion Date:</b> Jun-07				
<b>ASC nWBS Subprogram:</b> Facility Operations and User Support				
<b>Participating Sites:</b> LANL				
<b>Description:</b> Limited availability of Roadrunner capacity system for selected Stockpile applications. A number of Roadrunner Connected Units (CUs) will be integrated into the classified computing environment to support selected Stockpile applications. This integration includes the interconnect, the classified computing network, the global parallel file system, the HPSS archival storage system, the Network File System (NFS), and will include the system software environment required to deploy an initial Roadrunner system for limited availability.				
<b>Completion Criteria:</b> Ability to successfully run a selected application on the Roadrunner capacity system including access to all required classified network services (authentication, HPSS, etc.) including the PaScalBB global parallel file system.				
<b>Customer:</b> ASC and DSW Nuclear Weapons Programs are the Program Office level clients / customers for this milestone.				
<b>Milestone Certification Method:</b> Successful run of selected application within a pre-determined timeframe.				
<b>Supporting Resources:</b> SCC infrastructure, LANL classified network, HPSS.				
<b>Codes/Simulation Tools Employed:</b> TBD				
<b>Contribution to the ASC Program:</b> Meets compute cycle requirements for design codes.				
<b>Contribution to Stockpile Stewardship:</b> W76 and W88 milestones				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Selected Vendor delivering on time	High	Low	
2	Hiring for system integration	High	Medium	
3	Security Approval from DOE	High	Low	

<b>Milestone (ID#):</b> Develop and deploy Tripod capabilities for capacity computing environment (#2380)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY07
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Jun-07
<b>ASC nWBS Subprogram:</b> Facility Operations and User Support (FOUS) & Computational Systems and Software Environment (CSSE)
<b>Participating Sites:</b> LLNL, LANL, SNL
<p><b>Description:</b> The tri-labs will develop a coordination process and policies to deliver a seamless software environment for use by the NNSA tri-lab community (Los Alamos, Livermore and Sandia National laboratories), initially targeting Linux capacity computing clusters. The strategy has the following elements:</p> <ul style="list-style-type: none"> <li>• Identify current and near-term production quality solutions, and develop a tri-lab migration plan to this initial environment.</li> <li>• Identify tri-lab R&amp;D efforts focused on addressing gap areas, with acceptance criteria and an insertion plan for new technologies into the production stack.</li> <li>• Develop and implement tri-lab update decision, acceptance, and deployment policies and processes to maximize state-of-the-art software capabilities, while minimizing operational disturbance.</li> </ul> <p>The milestone deliverables are as follows:</p> <ol style="list-style-type: none"> <li><i>1. Deploy limited set of Tripod components on selected capacity systems at all three labs:</i> The tri labs will provide, at a minimum, selected applications development tools listed on the project website: <a href="http://tripod.ca.sandia.gov/">http://tripod.ca.sandia.gov/</a> under "ASC Project Milestones." Tasks required include: active participation in Tri-lab discussion of the essential common tools environment continues; Q3 deployment of current working version levels of the tool stack; appropriate documentation of user access methods for the tools and environment is made available to user; and procedures are in place for tracking, updating, and documenting version changes.</li> <li><i>2. Deploy selected tri-lab resource manager at LLNL and develop support model:</i> LLNL will integrate and support a commercial resource manager software product at LLNL to be used across the tri-lab HPC facilities. Tasks required include: 1) Integrate chosen tri-lab resource manager software product on at least one major computing resource at LLNL by end of Q2; and 2) develop a support plan that defines the roles and responsibilities between the tri-labs and the vendor for delivering a complete resource management solution. This plan will address the installation, configuration, and maintenance of the vendor's product; the vendor's responsibility to provide features that the tri-labs require; and the vendor's obligations to resolve problems encountered with the use of their product. This plan will also include a method for prioritizing the needs of the tri-lab facilities to ensure that the support the vendor provides is directed toward the greater good for the tri-labs. The support plan is to be drafted and approved by end of Q1. Los Alamos and Sandia will deploy this tri-lab resource manager in FY08.</li> <li><i>3. Produce a Tripod startup script at all three labs for prototyping:</i> Each lab will produce a startup script that could be used to put the tools that each site currently has licensed, and that are in the common toolset as per the Tripod tool list on</li> </ol>

<p>Tripod web site, into a user environment on at least one machine at each site for prototype purposes. An additional follow-on deliverable will be documentation on how to use that environment and a simple user test of each lab's Tripod environment. Sites may go beyond prototyping if resources become available.</p>				
<p><b>Completion Criteria:</b>                  1) Deployment of the capabilities listed above as the 3 milestone deliverables.                  2) Completion of the definition of the next set of Tripod capabilities, along with a tri-lab policies and process plan for deployment.</p>				
<p><b>Customer:</b> NNSA HQ</p>				
<p><b>Milestone Certification Method:</b> Completion evidence will be (1) documentation and hand-off of the FY07 Tripod deployment on the tri-labs' Linux capacity computing platforms; and (2) a report detailing the decision, acceptance, and deployment policies for the next-generation Tripod capabilities.</p>				
<p><b>Supporting Resources:</b></p>				
<p><b>Codes/Simulation Tools Employed:</b></p>				
<p><b>Contribution to the ASC Program:</b> Common environments leverage resources and enhance user efficiency.</p>				
<p><b>Contribution to Stockpile Stewardship:</b></p>				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure

## Milestone Descriptions for Preliminary FY08

<b>Milestone (ID#):</b> Deploy Tri-Lab Petascale Data Analysis Tool and Infrastructural Support (LL-CSSE-08-01)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q3
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment
<b>Description:</b>

<b>Milestone (ID#):</b> Uncertainty Quantification Application Demonstration of Design-Through-Analysis Environment (SNL-CSSE-08-01)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q4
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment
<b>Description:</b> Using an automated workflow model, which includes both pre- and post-processing components of the DTA environment we will demonstrate an uncertainty quantification capability using a specific weapon system application. The workflow will manage parameterized ensembles of simulations to evaluate the effect of important uncertainties for that application, which will be associated with one of the DART pilot projects. The automated workflow model will enable analysts to create and manage families of parameterized simulations to support uncertainty quantification. In general, model parameters can include geometric dimensions, mesh parameters, material properties, boundary and initial conditions, and other model attributes. The workflow will automatically and systematically select the appropriate parameter combination for each run in the ensemble and will manage the collective runs to generate a statistical analysis of the results. This work will involve integration of additional tools into the DTA environment, and will require interoperation with the Dakota framework and possibly commercial packages for additional statistical analysis of the ensemble simulation results. This work requires substantial advances in our ability to automate individual tool execution and to manage their interoperation.

<b>Milestone (ID#):</b> Tri-Lab Visualization Tool (LA-CSSE-08-01)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q3
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment
<b>Description:</b> Evaluation of Feasibility of tri-lab tool which unites and integrates the feature set for Paraview and Visit.

<b>Milestone (ID#):</b> Seamless and Productive Workflow (LA-CSSE-08-02)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q3
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment
<b>Description:</b>

<b>Milestone (ID#):</b> Demonstrate Heterogeneous Computing (LA-FOUS-08-01)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q2
<b>ASC nWBS Subprogram:</b> Facility Operations and User Support
<b>Description:</b>

<b>Milestone (ID#):</b> Infrastructure Plan for Petascale Computing Environment (TRI-CSSE-08-01)
<b>Level:</b> 2
<b>Fiscal Year:</b> FY08
<b>DOE Area/Campaign:</b> ASC
<b>Completion Date:</b> Q4
<b>ASC nWBS Subprogram:</b> Computational Systems and Software Environment
<b>Description:</b>

## VI. ASC Integration

In the tables below, ASC program management has described how the subprograms contribute to a milestone or another deliverable of any of the subprograms, products, or external programs for a given subprogram. Included in the descriptions is whether or not a delay would impact a planned deliverable or subprogram, product, or external program. If the table cell is blank, the product does not affect any deliverables.

**Table VI-1. All Subprograms as Related to Integrated Codes**

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
<b>Integrated Codes</b>	
LLNL	MMPC acts as the vehicle, which integrates the work from many other products to provide the complex weapons physics simulation tools. Legacy Codes act as a touchstone with the past and serve as a first check for the ability of MMPC to simulate UGTs. FRIC provides libraries (especially solvers) and other algorithms that are used by MMPC and Emerging and Specialized Codes. Emerging and Specialized Codes provide MMPC with nascent codes as they reach a sufficient level of maturity.
LANL	Theory and experiment developed in the Science Campaigns come together in the MMPC, which provide DSW with the simulation tools required for weapons certification, assessment, surety, SFI resolution, and design feasibility. Legacy codes are historical UGT calibration tools that provide a benchmark for the MMPC, with that benchmark being a rudimentary predictive simulation capability with uncertainties greater than those deemed acceptable by the MMPC. FIRC Provides for higher-risk, exploratory research (Weapons-Supported Research) necessary to pursue high-payoff numerical algorithm innovations. Deliverables include more robust and accurate algorithms for MMPC in the area of shock physics, transport, linear and nonlinear solvers, and coupled multi-physics. Provides an analysis tool for the design of casting /welding operations in the nuclear weapons complex, setup tools for providing mesh geometry and initial and boundary condition information for the MMPC, and energy diagnostic tools for post-processing the MMPC output in weapons effects studies. The casting simulation tool contributes reusable software components (solvers, parallel communication, etc.) for all of IC Products.

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
SNL	FRIC Provides the libraries and frameworks, including solvers, optimization other algorithms that are used by integrated codes at SNL, particularly engineering codes.
<b>Physics and Engineering Models</b>	
LLNL	Improved models of physical behavior—including, strength, equation of state, nuclear reactions, opacity, high explosive performance, and advanced hydrodynamic behavior—are delivered for use in the modern multi-physics codes, leading towards the goal of improved predictive capability using science based models. Contributes to FY06 L1 and L2 Integrated Codes milestones, and a proposed FY07 L2 milestone. Delays in this product impact the progress in improving the predictive capability of the integrated codes. Model Implementation facilitates the delivery of an effective capability within the modern multi-physics codes, moving forward with improved predictivity. New model implementations are achieved through close collaboration between the integrated code developers and the developers of the new physics models. In some cases the ASC integrated codes are utilized as the development test-bed for new models. Contributes to FY06 L1 and L2 Integrated Codes milestones, and a proposed FY07 L2 milestone. Delays in this product impact the progress in improving the predictive capability of the integrated codes. This activity provides new material and physics models for integration in the MMPC and possibly ESC and legacy codes. Fundamental Physics Codes and Application delivers codes that provide opacity data that is directly inserted into integrated code simulations. Contributes to FY06 and FY07 L2 Integrated Codes milestones. Delays in this product impact the progress in improving the predictive capability of the integrated codes through accuracy of simulations and increased uncertainties. Material Data Libraries delivers updated libraries of properties of programmatic materials for various physics packages, including equation of state, opacity, and nuclear reactions, to the integrated codes. Library maintainers work closely with the integrated code developers to insure effective data access through the library APIs. Contributes to FY06 L1 and L2 Integrated Codes milestones. Delays in this product impact the accuracy of simulated predictions, increasing uncertainty, and delaying DSW applications. This activity provides new material data libraries for integration in the MMPC and possibly ESC and legacy codes.
LANL	Theoretical Models and Experimental Integration Provides for development of fundamental, first-principles theory in area of nuclear and atomic physics, equations of state, high explosives, materials response and damage, turbulence and mix, thermonuclear burn, and engineering analysis. Also provides for higher-risk, exploratory research (Weapons-Supported Research) necessary to pursue high-payoff weapons physics innovations. These theoretical developments drive the improved models instantiated in the Model Implementation Products that are delivered into the MMPC. Implements (in software libraries) new and improved models developed in the Theoretical Models and Experimental Integration Product for the MMPC and legacy codes in area of materials response and damage, equations of state, turbulence and mix, thermonuclear burn, and engineering analysis. Fundamental Physics Codes and Application

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
	Develops, maintains, evolves, and applies codes whose algorithms represent high-fidelity solutions to fundamental, first-principles formulations for nuclear and atomic physics, equations of state, materials response, high explosives, turbulence and mix, and thermonuclear burn. These codes are the vehicle by which new, postulated theories developed in the Theoretical Models and Experimental Integration Product are verified and validated. Successfully validated theories are delivered as new models and materials libraries to the MMPC. Material Data Libraries Integrates and tests software libraries and databases into the Integrated Codes Sub-Program (primarily the MMPC product) that instantiate nuclear and atomic physics data, equations of state, and high explosive and foam materials response.
SNL	Model implementation provides science-based sub-grid and constitutive models for the ASC Engineering Codes. This activity requires tight integration between project leads and code PIs.
<b>Verification and Validation</b>	
LLNL	Establishes the validity of the models and code functionality required to assess Performance. May confirm the adequacy, or otherwise, of developments in the MMPC, ESC or Legacy codes and hence may provide new requirements. V&V Methods supports all aspects of software quality assurance for the Integrated Codes as well as the development of mathematical methods for assessing Uncertainty Quantification of the outputs of the Integrated Codes. Specialized V&V Assessments are used to make targeted assessment of code performance against specific stockpile weapons systems or families of systems. Data Validation and Archiving is used to support the variety of databases that the Integrated Codes access to perform their function. This activity may provide improved experimental data with which to aid the validation of models in MMPC, ESC or legacy codes.
LANL	Develops fundamental V&V and QMU methodologies and test suites to be used in periodic formal and quantitative assessments of codes primarily released in the MMPC Product area. The test suites are archived, assessments are documented, and methodologies are delivered as standard tools. Also supports software quality assurance activities for all codes developed in the Integrated Codes Sub-Program. Develops and conducts, for performance codes in the MMPC Product area, test suites, sensitivity analyses, uncertainty quantification, single- and multi-effects physics validation, and verification studies. These assessments, which are conducted and documented regularly, are part of an evolutionary delivery life cycle in which primary performance simulation requirements are progressively elaborated. Develops and conducts, for selected SNL and LLNL engineering codes delivered in the Integrated Codes Sub-Program, test suites, sensitivity analyses, uncertainty quantification, single- and multi-effects physics validation, and verification studies. These assessments, which are conducted and documented regularly, are part of an evolutionary delivery life cycle in which engineering simulation requirements are progressively elaborated. Performs focused V&V assessments of specific physics or combinations of physics of interest to the Integrated Codes Sub-Program, such as astrophysical applications

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
	(for radiation hydrodynamics) and output diagnostics (for transport). Establish requirements and deliver the software and hardware infrastructure for maintaining archival databases needed to support V&V assessments of codes in the Integrated Codes Sub-Program. These databases include (but are not limited to) data generated in small-scale and integral experiments as well as UGT events.
SNL	The Engineering Codes and Legacy Code CTH rely on verification methodologies developed in the V&V Methods product to execute software quality practices. Methodologies for verification, validation, UQ and Sensitivity analyses are used extensively through V&V applications and all L1 and L2 milestones to which the V&V program contributes. Directly supports L1 milestones and supporting L2 milestones; Directly supports Alegra-HEDP Z-Pinch milestone.
<b>Computational Systems and Software Environment</b>	
LLNL	Contributes the required computational resources to enable the development, test, and usage of the 3D and full systems physics for the large codes. A delay in delivery of capability systems impacts the resources available to the entire tri-lab program, making it more difficult to provide the Integrated code capabilities. Contributes the required computational resources for code development and production parameter studies using the products of the Integrated Codes. Delay or inadequate capacity system resources slow down the progress of this area. Contributes to the resources needed for research in preparing for future petascale calculations needed by the program in about 2010. Contributes the software environment needed to operate the computational resources for the development, test, and production usage of the new codes. A delay in delivery of system software and tools impacts the resources available to the entire tri-lab program, making it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted. Contributes the files systems and storage systems needed to operate the computational resources for the development, test, and production usage of the codes. A disruption or delay in delivery of these systems may make the computational resources needed by Integrated codes unusable. A disruption or delay of the networking developments may impact the remote usage of the capability systems by the Integrated Code efforts at the other laboratories. The Post Processing product contributes to the understanding of the results of the Integrated Codes. For the large codes on the large systems, the volume of data produced is frequently too great to make judgment about the quality of the result without visual processing of the results. Delay in deliverables from this product impacts productivity of the Integrated Code efforts.
LANL	Ensures the hardware and software maintenance, reliability, availability, and security of LANL capability platforms (currently the Q platform) for reliable execution of MMPC codes during the simulation of high-fidelity, 3D, full physics scenarios. Procures, integrates, and maintains the capacity computing resources required for development of codes in the Integrated Codes Sub-Program and their end use application for DSW, certification, assessment, and SFI resolution.

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
	Perform R&D work with selected computer vendors in the design and demonstration of next-generation (>3-5 years away) computing platforms whose characteristics are tuned for more efficient execution performance of codes in the MMPC Product area. Provide a reliable and resilient system software environment for codes in the Integrated Codes Sub-Program by focusing on Linux software stack extensions (Science Appliance), efficient and portable MPI, useful IDEs (Eclipse environment), code coverage tools (Alexandria), process migration, scalable run-time, and network device failover. Understand, predict, and improve performance of integrated codes on these computer systems through performance modeling and analysis. Provide, support, and evolve the hardware and software necessary to meet the run-time and archival output requirements for codes in the Integrated Codes Sub-Program.
SNL	Robust operating systems are required for Engineering Codes and Legacy Codes to run reliably on capability platforms. Robust operating systems are required for Engineering Codes and Legacy Codes to run reliably on capacity platforms. Sufficient computing capacity is required to meet the user demand for simulations with Engineering Codes and Legacy Codes. Provides the computing systems focused on the mid-range capability gap (workloads in the range of approximately 500-4,000 processors). Supports both production computing workload and application/algorithm development resources. May also provide resources for jobs below 500 processors. Robust operating systems that run predictably on various platforms are required for V&V assessments and applications requiring capacity computing. Sufficient user capacity required to meet needs of V&V and applications, especially those requiring QMU. Provides the system software for both Capability and Capacity Systems the will be used by Integrated Codes workload. This Product also provides the code development tools to support both Application and Algorithm development. Provides and supports critical infrastructure required for code development and production simulations. In FY06, this product is completing the production deployment of file systems, storage and networking for the Red Storm general availability computing environment; developing the next generation security infrastructure in support of tri-lab computing; and laying the groundwork in the form of a roadmap to deliver petaflops computing. SNL: Fast creation of meshes and input files for Engineering Codes and Legacy Codes are required for efficient and effective simulations with these codes. Visualization capabilities on both the desktop and other platforms are needed to interpret results from Engineering Codes and Legacy Codes.
<b>Facility Operations and User Support</b>	
LLNL	This product is required for all production operation of the computational resources needed by this subprogram, including tri-lab system access. Contributes the user support needed to access and operate the computational resources for the development, test, and production usage of the integrated codes. A disruption or delay in delivery of user support impacts the resources available to the entire tri-lab program, making it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted.

	<b>Integrated Codes (Modern, Multi Physics Codes; Legacy Codes; Engineering Codes; Focused Research, Innovation, and Collaboration; and Emerging and Specialized Codes)</b>
LANL	Deliver systems and network operations and maintenance and facilities maintenance and upgrades for all capacity and capability computing platforms used by codes in the Integrated Codes Sub-Program. Provide on-call experts for timely consultation with end users and developers of codes in the Integrated Codes Sub-Program experiencing software and hardware problems on all computing platforms.

**Table VI-2. All Subprograms as Related to Physics and Engineering Models**

	<b>Physics and Engineering Models (Theoretical Models and Experimental Integration; Model Implementation; Fundamental Physics Codes and Application; Material Data Libraries)</b>
<b>Integrated Codes</b>	
LLNL	Supplies requirements and feedback to PEM regarding Material Data Libraries and models.
LANL	Supply requirements and feedback to the PEM Sub-Program through (integral, regression, unit) testing and end use on the accuracy and robustness of delivered and integrated material data libraries and models for material response, turbulence, burn, and high explosive behavior. Supply requirements and feedback to the PEM Sub-Program through (integral, regression, unit) testing on the impacts of new material data libraries and models for material response, turbulence, burn, and high explosive behavior to algorithm changes in the areas of hydro, transport, solvers, and multi-physics coupling.
SNL	In some case changes to the codes may be required to handle advanced constitutive models or sub-grid models. The PEM staff and management work with PIs from Engineering codes.
<b>Physics and Engineering Models</b>	
LLNL	This product provides sub-continuum scale simulations, which are used in the development of new theoretical models as well as in producing parameterizations of physical properties used to populate material data libraries.

<b>Physics and Engineering Models (Theoretical Models and Experimental Integration; Model Implementation; Fundamental Physics Codes and Application; Material Data Libraries)</b>	
LANL	Provides for development of fundamental, first-principles theory in area of nuclear and atomic physics, equations of state, high explosives, materials response and damage, turbulence and mix, thermonuclear burn, and engineering analysis. Also provides for higher-risk, exploratory research (Weapons-Supported Research) necessary to pursue high-payoff weapons physics innovations. These theoretical developments drive the improved models instantiated in the Model Implementation Products. Implements (in software libraries) new and improved models developed in the Theoretical Models and Experimental Integration Product in area of materials response and damage, equations of state, turbulence and mix, thermonuclear burn, and engineering analysis. Develops, maintains, evolves, and applies codes whose algorithms represent high-fidelity solutions to fundamental, first-principles formulations for nuclear and atomic physics, equations of state, materials response, high explosives, turbulence and mix, and thermonuclear burn. These codes are the vehicle by which new, postulated theories developed in the Theoretical Models and Experimental Integration Product are verified and validated.
SNL	Contributes to the model implementation product. Gives theoretical underpinnings for sub-grid and constitutive models. Works with Model Implementation PIs. Model implementation is dependent on successful theoretical activities under the above product. These products are tightly integrated.
<b>Verification and Validation</b>	
LANL	Develops fundamental V&V and QMU methodologies and undertakes software quality assurance activities that are used in the assessment of material data libraries and models delivered for use in codes in the Integrated Codes Sub-Program. Scrutinizes the material data libraries and models delivered by PEM to performance codes by conducting sensitivity analyses, uncertainty quantification, single- and multi-effects physics validation, and verification studies. Scrutinizes material (nuclear and atomic) data libraries delivered by PEM to astrophysical applications (for radiation hydrodynamics) and output diagnostics (for transport). Establish and maintain experimental databases needed to validate new models and material data libraries developed by PEM.
SNL	Verification and validation of sub-grid and constitutive models is reliant on V&V methods and engineering assessments to provide validated models.

<b>Physics and Engineering Models (Theoretical Models and Experimental Integration; Model Implementation; Fundamental Physics Codes and Application; Material Data Libraries)</b>	
<b>Computational Systems and Software Environment</b>	
LLNL	<p>A delay in delivery of capability systems impacts the total resources available to the entire tri-lab program, limiting resources available to the Physics and Engineering model development. Contributes the required computational resources for code development and production parameter studies using the products of the Physics and Engineering models. Delay or inadequate capacity system resources slows down the progress of this area. The BlueGene/L system will be directly used for materials aging studies. Contributes the software environment needed to operate the computational resources for the development, test, and production usage of the new models and codes – especially for the BlueGene/L environment. A delay in delivery of system software and tools support makes it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted. Contributes the files systems and storage systems needed to operate the computational resources for all code use. A disruption or delay in delivery of these systems may make the computational resources unusable. The function of the parallel files systems will be especially important for the effective use of BlueGene/L for material studies. The Pre- and Post-Processing product contributes to an understanding of the Physics and Engineering Models. For highly scaled physics and materials codes on BlueGene/L, the volume of data produced coupled with unique architecture of BlueGene/L requires next-generation data visualization and analysis tools and servers. Delay in tool and server deployment from this product will have a large impact on the productivity and understanding of the Physics and Engineering Models efforts.</p>
LANL	<p>Ensures the hardware and software maintenance, reliability, availability, and security of LANL capability platforms (currently the Q platform) with which to scrutinize the effects of PEM material data libraries and models on predictive simulation of UGTs and full system scenarios. Procures, integrates, and maintains the capacity computing resources required for development and testing of PEM material data libraries and models and their integration into codes in the Integrated Codes Sub-Program. Perform R&amp;D work with selected computer vendors in the design and demonstration of next-generation (&gt;3-5 years away) computing platforms that are cognizant of the memory bandwidth requirements imposed by the material data libraries and models supplied by PEM. Provide a reliable and resilient system software environment for material library and model development in the PEM Sub-Program by focusing on Linux software stack extensions (Science Appliance), efficient and portable MPI, useful IDEs (Eclipse environment), coverage tools (Alexandria), process migration, scalable run-time, and network device failover. Provide, support, and evolve the hardware and software necessary to meet the run-time and archival output requirements for material data libraries and models implemented in the PEM Sub-Program.</p>

	<b>Physics and Engineering Models (Theoretical Models and Experimental Integration; Model Implementation; Fundamental Physics Codes and Application; Material Data Libraries)</b>
SNL	Provides the computing systems that are focused on the Mid-range Capability Gap, such as, workloads in the range of approximately 500-4,000 processors. Supports both production computing workload and Application and Algorithm development resources. May be well suited for extended job duration workloads.
<b>Facility Operations and User Support</b>	
LLNL	This product is required for all production operation of the computational resources needed by this subprogram, including tri-lab system access. Contributes the user support needed to access and operate the computational resources for the development, test, and production usage of the codes and applications. A disruption or delay in delivery of user support impacts the resources available to the entire tri-lab program, making it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted.
LANL	Deliver systems and network operations and maintenance and facilities maintenance and upgrades for all capacity and capability computing platforms used by the PEM Sub-Program. Provide on-call experts for timely consultation with developers and integrators of models and material data libraries in the PEM Sub-Program experiencing software and hardware problems on all computing platforms.

**Table VI-3. All Subprograms as Related to Verification and Validation**

	<b>Verification and Validation (V&amp;V Methods, Primary V&amp;V Assessments, Secondary V&amp;V Assessments, Engineering V&amp;V Assessments, Specialized V&amp;V Assessments, Data Validation and Archiving, University Partnerships)</b>
<b>Integrated Codes</b>	
LLNL	MMPC provides the code capabilities for V&V assessments.
LANL	Provides the codes used in primary and secondary V&V assessments, acts on corrective actions resulting from those assessments, and improves software quality as guided by activities and training in the V&V Sub-Program. Improves software quality as guided by activities and training in the V&V Sub-Program. Uses findings in primary and secondary and V&V assessments to provides more robust and accurate (shock physics, transport, linear and nonlinear solvers, and coupled multi-physics) algorithms. Improves software quality as guided by activities and training in the V&V Sub-Program. Improves software quality as guided by activities and training in the V&V Sub-Program and acts on corrective

	<b>Verification and Validation (V&amp;V Methods, Primary V&amp;V Assessments, Secondary V&amp;V Assessments, Engineering V&amp;V Assessments, Specialized V&amp;V Assessments, Data Validation and Archiving, University Partnerships)</b>
	actions identified in V&V assessments.
SNL	The Engineering Codes CALORE, FUEGO, SALINAS and DAKOTA will be used in executing the FY06 tri-lab level 1 V&V milestone. Normal user support is expected to be required, but additional support will be supplied if required. The code Alegra-HEDP will support a L2 V&V milestone. V&V and QMU applications require that codes are ready for stockpile computing such as, there is adequate SQE, regression testing, code verification, and acceptance tests. Provides the optimization framework and algorithms that are used by the verification and validation products at SNL. Provides the framework that is used by SNL, LLNL and LANL verification and validation products. The V&V program and QMU studies will rely on capabilities afforded by Dakota to perform UQ studies, sensitivity studies, and optimization studies.
<b>Physics and Engineering Models</b>	
LLNL	Provides guidance on experiments necessary for verifying new physics and engineering models. Provides the most up to date material property libraries for use in the integrated codes for doing the various V&V system assessments. Provides the best available physics models within the integrated codes for doing the various V&V system assessments.
LANL	Communicates verification and validation requirements imposed by new and improved theoretical models and works with V&V assessment teams in developing metrics, test suites, etc. Ensures that the latest software libraries embodying new and improved models in area of materials response and damage, equations of state, turbulence and mix, thermonuclear burn, and engineering analysis are available to V&V assessment teams. Ensures that the latest software libraries and databases embodying nuclear and atomic physics data, equations of state, and high explosive and foam materials response are available to V&V assessment teams. Improves software quality in fundamental codes as guided by activities and training in the V&V Sub-Program. Communicates validation requirements imposed by the theoretical models embodied in these codes to V&V assessment teams.
SNL	Models from PEM will be required for L1 milestone and support L2 milestones. Models from PEM are required for other V&V assessments. PEM models must be verified and validated in a manner consistent with the requirements of the V&V program and applications that require QMU assessments. V&V program requires a consistency and pedigree of material data and its usage.
<b>Verification and Validation</b>	
LLNL	Develops fundamental V&V and QMU methodologies and test suites to be used in periodic formal and quantitative code assessments by the primary, secondary, and engineering assessment teams. The test suites are archived, assessments are documented, and methodologies are delivered as standard tools.

	<b>Verification and Validation (V&amp;V Methods, Primary V&amp;V Assessments, Secondary V&amp;V Assessments, Engineering V&amp;V Assessments, Specialized V&amp;V Assessments, Data Validation and Archiving, University Partnerships)</b>
LANL	Conducts code V&V assessments that utilize the methodologies and tools devised by the V&V Methods Product area. Work with other assessment teams in evolving the assessment rigor, formality, and corrective action feedback. Support V&V assessments by delivering the software and hardware infrastructure necessary for maintaining archival databases of data generated in small-scale and integral experiments as well as UGT events.
SNL	Methodologies for verification, validation, UQ and Sensitivity analyses are used extensively through V&V applications and all L1 and L2 milestones to which the V&V program contributes. Directly supports L1 milestones and supporting L2 milestones; Directly supports Alegra-HEDP Z-Pinch milestone.
<b>Computational Systems and Software Environment</b>	
LLNL	Contributes the required computational resources to enable the test, and usage of the 3D and full systems physics for the large codes. A delay in delivery of capability systems impacts the resources available to the entire tri-lab program, making it more difficult to perform the V&V testing of the Integrated code capabilities. Contributes the required computational resources for production parameter studies needed for V&V. Delay or inadequate capacity system resources slows down the progress of this area. Research and experience in advanced systems will contribute to preparation for future petascale calculations needed for future stockpile calculations. Contributes the software environment needed to operate the computational resources for the production V&V testing. If the environment isn't robust, it is more difficult to use the computational resources effectively and user productivity / time-to-solution is impacted. Contributes the files systems and storage systems needed to operate the computational resources for all production use. V&V activities also rely on access to older archived results. The Pre- and Post-Processing product helps support new computational paradigms and increased program emphasis of V&V activities. The large number of runs (100s to 1000s) and vast amount of (possibly geographically distributed) data requires new ways to manage and analyze large data. Delay in this product's deployment will have a large impact on the ability of V&V to analyze and post-process results in a timely way.
LANL	Ensures the hardware and software maintenance, reliability, availability, and security of LANL capability platforms (currently the Q platform) with which to conduct V&V assessments of predictive simulation of UGTs and full system scenarios. Procures, integrates, and maintains the capacity computing resources required for development of V&V Methods products and conducting the primary, secondary, and engineering V&V code assessments. Performs R&D work with selected computer vendors in the design and demonstration of next-generation (>3-5 years away) computing platforms that are cognizant of the computational loads imposed by rigorous V&V assessments, and UQ/sensitivity studies. Provides a reliable and resilient system software environment for code assessments in the V&V Sub-Program by focusing on Linux software stack extensions (Science Appliance), efficient and portable MPI, useful IDEs (Eclipse

	<b>Verification and Validation (V&amp;V Methods, Primary V&amp;V Assessments, Secondary V&amp;V Assessments, Engineering V&amp;V Assessments, Specialized V&amp;V Assessments, Data Validation and Archiving, University Partnerships)</b>
	environment), coverage tools (Alexandria), process migration, scalable run-time, and network device failover. Provide, support, and evolve the hardware and software necessary to meet the run-time and archival output requirements for V&V assessments in the V&V Sub-Program. Provide for current and next generation production visualization needs of V&V assessments in the V&V Sub-Program. Associated activities include visualization hardware and user support of existing systems, next generation visualization R&D, expert visualization, legacy visualization tool support, and quantitative data analysis
SNL	Provides the computing systems that are focused on the Mid-range Capability Gap, such as, workloads in the range of approximately 500-4,000 processors. Supports both production computing workload and V&V and QMU Algorithm development resources. May be especially useful for parameter studies and QMU assessments. Provides the system software for both Capability and Capacity Systems the will be used by V&V and QMU workload.
<b>Facility Operations and User Support</b>	
LLNL	This product is required for all production operation of the computational resources needed by this subprogram, including tri-lab system access. Contributes the user support needed to access and operate the computational resources for data validation and archiving, and University partnerships. A disruption or delay in delivery of user support impacts the resources available to the entire tri-lab program, making it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted.
LANL	Deliver systems and network operations and maintenance and facilities maintenance and upgrades for all capacity and capability computing platforms used by the V&V Sub-Program. Provide on-call experts for timely consultation with V&V assessment teams in the V&V Sub-Program experiencing software and hardware problems on all computing platforms.
SNL	The V&V program has strong collaborations with the University of Utah in fire modeling and V&V of fire modeling codes; The V&V program has a collaboration with Stanford on turbulent fluid dynamics.

**Table VI-4. All Subprograms as Related to Computational Systems and Software Environment**

<b>Computational Systems and Software Environment (Capability Systems; Capacity Systems; Advanced Systems, System Software and Tools; I/O, Storage Systems; and Networking; Pre- and Post-Processing Environments)</b>	
<b>Integrated Codes</b>	
LLNL	Supplies requirements and feedback to CSSE to improve systems and environments deployed for ASC.
LANL	Supplies requirements and feedback to the CSSE Sub-Program on capability, capacity, and advanced systems hardware, software, and environments through code team development, alpha use, and testing activities. Work with the CSSE code optimization and performance modeling efforts in identifying and improving code execution performance.
SNL	There are requirements for the Engineering Codes to be running on ASC capability and capacity platforms. Provides the fundamental research in algorithms, advanced architectures, performance modeling, I/O and pre and post processing.
<b>Physics and Engineering Models</b>	
LLNL	Works closely with the staging teams for new capability and capacity systems; frequently the science codes are the early adopters of new systems and are used during the testing period to provide a thorough shake-out of new machines. Contributes to certifying that new systems are ready for general availability.
LANL	Supply requirements and feedback to the CSSE Sub-Program on capability and capacity systems hardware, software, and environments used to support theoretical model development, in implementing models into codes in the Integrated Codes Sub-Program, in developing and applying fundamental physics codes, and in developing software libraries and databases that instantiate nuclear and atomic physics data, equations of state, and high explosive and foam materials response Work with the CSSE code optimization effort in identifying and improving execution performance.
<b>Verification and Validation</b>	
LLNL	This product contributes requirements and feedback to the CSSE subprogram.
LANL	Supply requirements and feedback to the CSSE Sub-Program on capability and capacity systems hardware, software, and environments used to develop tools in the V&V methods, code assessment, and data archiving activities.

	<b>Computational Systems and Software Environment (Capability Systems; Capacity Systems; Advanced Systems, System Software and Tools; I/O, Storage Systems; and Networking; Pre- and Post-Processing Environments)</b>
<b>Computational Systems and Software Environment</b>	
LLNL	A delay in the deployment of the Purple system and BlueGene/L will result in a direct delay of other milestones in this subprogram. The capacity systems supplied by this product will be the development target of other products in this subprogram. If the software environment is not deployed in a timely fashion, or is not robust – other milestones and efforts in this subprogram are directly impacted and delayed; likewise if the file systems are not operational and robust – other milestones and efforts in this subprogram are directly impacted and delayed. The Pre- and Post-Processing product delivers both hardware and software for the Computational Systems and Software Environment. Delays in researching, developing, integrating and deploying new data analysis and visualization tools will impact the ability of this subprogram to meet its overall goals and milestones.
LANL	Leverage, where possible, any hardware and software maintenance, reliability, availability, and security activity in support of LANL capability platforms (currently the Q platform) to similar capacity platform activities. Share computing facility requirements with other (capacity, advanced) systems. Comply with requirements and resources needed for software development activities in the CSSE Sub-Program. Transfer any mature technologies developed as a result of R&D work with selected computer vendors in the design and demonstration of next-generation (>3-5 years away) computing platforms to appropriate production efforts in the CSSE Sub-Program. Provide a reliable and resilient capability and capacity system software environment by focusing on Linux software stack extensions (Science Appliance), efficient and portable MPI, useful IDEs (Eclipse environment), code coverage tools (Alexandria), process migration, scalable run-time, and network device failover. Communicate problems, requirements, and corrective actions to other Product areas in CSSE. Provide, support, and evolve the hardware and software on capability and capacity systems necessary for application run-time and archival output needs. Provide feedback and requirements for capability and capacity system software and hardware necessary for current and next generation production visualization of application output data.
SNL	Provides development resources and some opportunity for large-scale Capacity dedicated system time testing and evaluation. Will support integration and deployment of products from other CSSE product areas. Provides the runtime, system software and system integration and management tools for Capability, Capacity and Advanced Systems the will be procured by CSUE product areas. This Product also provides the application performance and RAS Metrics that can be useful to guide decisions to upgrade or change system software. System integration and management tools are also used by many Pre and Post processing systems.

	<b>Computational Systems and Software Environment (Capability Systems; Capacity Systems; Advanced Systems, System Software and Tools; I/O, Storage Systems; and Networking; Pre- and Post-Processing Environments)</b>
<b>Facility Operations and User Support</b>	
LLNL	This product is required for the operational aspects needed to support the work of all of the efforts in this subprogram. User Support services provides the direct source of feedback from customers about problems related to this CSUE subprogram. It also provides the documentation and training need by the subprogram products. A disruption or delay in delivery of user support impacts the resources available to the entire tri-lab program, making it more difficult to use the computational resources effectively. User productivity and time-to-solution is impacted. This product includes procurement support for acquisitions of this subprogram.
LANL	Deliver systems and network operations and maintenance and facilities maintenance and upgrades for all capacity and capability computing platforms. Provide on-call experts for timely consultation with developers of system software in the CSSE Sub-Program experiencing software and hardware problems on all computing platforms.

**Table VI-5. All Subprograms as Related to Facility Operations and User Support**

	<b>Facility Operations and User Support (Facilities, Operations, and Communications; User Support Services; Collaborations)</b>
<b>Integrated Codes</b>	
LLNL	Supplies requirements and feedback to FOUS to improve facilities, operations, and support services.
LANL	Supply requirements and feedback to the FOUS Sub-Program on the facilities, operations, and support of capability and capacity systems through code team development, alpha use, and testing activities.
SNL	Provides an infrastructure for external collaborations, including facilities, administration and networking.
<b>Physics and Engineering Models</b>	
LLNL	Contributes requirements and feedback to the general operations and support of the computer facilities.
LANL	Supply requirements and feedback to the FOUS Sub-Program on the facilities, operations, and support of capacity systems used to support theoretical model development, in implementing models into codes in the Integrated Codes

	<b>Facility Operations and User Support (Facilities, Operations, and Communications; User Support Services; Collaborations)</b>
	Sub-Program, in developing and applying fundamental physics codes, and in developing software libraries and databases that instantiate nuclear and atomic physics data, equations of state, and high explosive and foam materials response.
<b>Verification and Validation</b>	
LLNL	Contributes requirements and feedback to the general operations and support of the computer facilities.
LANL	Supply requirements and feedback to the FOUS Sub-Program on the facilities, operations, and support of capability and capacity systems used in V&V code assessment and storage hardware used in data archiving activities.
<b>Computational Systems and Software Environment</b>	
LLNL	A delay in the deployment of the Purple and BlueGene/L systems will impact delivery of other resources and milestones supplied by this subprogram. The capacity systems supplied by this product will be the development target of other efforts in this subprogram. Problems with their operational status will be the direct responsibility of this subprogram. If the software environment, and archive and file systems, are not operational robust – a substantial load is put on the operational staff and user services staff to track problems and maintain operational status. ( Pre- and Post-Processing) Delivers both hardware and software for the Computational Systems and Software Environment. Delays in the procurement, installation, testing and deployment of data visualization and analysis hardware and tools will impact the ability of this subprogram to meet its overall goals and milestones.
LANL	Share computing facility operations and support requirements with other systems. Supply requirements and feedback to the FOUS Sub-Program on the facilities, operations, and support of capability and capacity systems used in system software development activities, for application run-time and archival output needs, and in current and next generation production visualization of application output data.
SNL	Provides the computing systems that are focused on the Mid-range Capability Gap. Close coordination and integration with 1.5.5 will be required to ensure Capacity System infrastructure requirements are met, such as, space, power, cooling, and network interfaces to storage, visualization and archiving resources. New capacity systems will also require staffing from 1.5.5 for system administration and user support.
<b>Facility Operations and User Support</b>	
LLNL	Interruptions in the facilities supported by this product causes a major load on the User Support services to interact with customers. User support services provides the direct customer problem tracking related to the RAS of the operational systems and tri-lab communications networks. Includes procurement support for acquisitions of this

	<b>Facility Operations and User Support (Facilities, Operations, and Communications; User Support Services; Collaborations)</b>
	subprogram.
LANL	Use experiences gained in facilities operations, maintenance, and upgrades to anticipate user problems areas, thereby coordinating fast and effective user support services. Track issues brought up by users and devise corrective actions for those issues directly related to facility operations and maintenance.

**Table VI-6. All Subprograms as Related to External Programs**

	<b>External Programs</b>
<b>Integrated Codes</b>	
LLNL	Simulation codes for Nuclear Explosive Package (NEP) engineering support DSW and LEP work at both LLNL (for example, W80-3) and LANL (for example, W76-1).
LANL	Provide DSW and the Science Campaigns with the simulation tools required for weapons certification, assessment, surety, SFI resolution, and design feasibility. These same tools are also feasible for design and assessment simulations of DOD conventional weapons. Provide DSW with the historical UGT calibration simulation tools required for benchmarking and quantifying simulation uncertainties (relative to the MMPC codes) in weapons certification, assessment, surety, SFI resolution, and design feasibility. Targets the reduction of numerical solution uncertainties known to exist in DSW certification by engaging for higher-risk, exploratory research (Weapons-Supported Research) necessary to pursue high-payoff numerical algorithm innovations in the area of shock physics, transport, linear and nonlinear solvers, and coupled multi-physics. Deliver simulation analysis and design support for W88 pit manufacturing, DOE Complex (e.g., LANL Sigma Foundry) casting design operations, Y-12 casting and welding operations, DOE Complex criticality safety scenarios, and DSW diagnostics output assessments and predictions.
<b>Physics and Engineering Models</b>	
LANL	Deliver theories useful in understanding basic physical phenomena of interest to experimental programs in DSW, the Science Campaigns, the DOD, and modeling and simulation programs in the DOE Office of Science. Software libraries embodying models for phenomena like materials response and damage and high explosive behavior can be

<b>External Programs</b>	
	useful for experimental programs in DSW, the Science Campaigns, for nuclear reactor programs in the NNSA, and for conventional weapons modeling and simulation program in the DOD. Atomistic, molecular dynamic, and chemistry codes generated in this product area represent tools for studying fundamental tools for studying and understanding phenomena such as material properties and behavior. These tools are need in many other programs such as those in the DOE Office of Science.
<b>Verification and Validation</b>	
LANL	Generic code V&V assessment tools (error estimates, QMU, convergence analysis), processes, and methodologies can be applied to similar assessments of simulation tools in other programs such as those being developed in the DOD and DOE Office of Science. These assessments lead to corrective actions that ultimately improve the code predictability, hence improve the simulation fidelity used for DSW certification. Generic code V&V assessment processes and methodologies can be applied to similar assessments of simulation tools in other programs such as those being developed in the DOD and DOE Office of Science. Archival databases of UGT events and other experiments is useful for DSW activities.
<b>Computational Systems and Software Environment</b>	
LANL	Other programs in need of and using HPC can benefit from ASC experiences and expertise gained in fielding capacity systems. Specifically, selected system software (e.g., Linux software stack) and vendor hardware used and developed for ASC platforms helps to advance future systems procured by other programs. Accelerates and potentially redirects vendor HPC strategies and plans to be more cognizant of HPC needs in scientific computing. This benefits other programs in need of and using HPC. System software and tools developed by ASC can ultimately find their way into commercial systems or standard operating systems (e.g., Linux). I/O and storage R&D driven by ASC capability and capacity needs helps to accelerate the availability and reliability of commercial vendors in this area (e.g., Panasas, Lustre, etc.). Visualization advances and needs helps to accelerate the availability, reliability, and usefulness of commercial vendor products (e.g., Computational Engineering International's EnSight) in this area.
<b>Facility Operations and User Support</b>	
LANL	Knowledge and experience gained from operations, maintenance, and upgrades of HPC facilities is useful for other programs engaged in HPC activities, such as Program Offices in the DOD and DOE Office of Science.

## VII. ASC Risk Management

Risk management is a process for identifying and analyzing risks, executing mitigation and contingency planning to minimize potential consequences of identified risks, and monitoring and communicating up-to-date information about risk issues. Risk management is about identifying opportunities and avoiding losses. A “risk” is defined as (1) a future event, action, or condition that might prevent the successful execution of strategies or achievement of technical or business objectives, and (2) the risk exposure level, defined by the likelihood or probability that an event, action, or condition will occur, and the consequences, if that event, action, or condition does occur. Table VII-1 summarizes ASC’s top ten risks, which are managed and tracked.

**Table VII-1. ASC’s Top Ten Risks<sup>7</sup>**

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
1	Compute resources are insufficient to meet capacity and capability needs of designers, analysts, DSW, or other Campaigns.	High	High	HIGH	Integrate program planning with DSW and other Campaigns, to ensure requirements for computing are understood and appropriately set; maintain emphasis on platform strategy as a central element of the program; pursue plans for additional and cost-effective capacity platforms.
2	Designers, analysts, DSW, or other Campaign programs lack confidence in ASC codes or models for application to certification /qualification.	Very High	Low	MEDIUM	Maintain program emphasis on V&V; Integrate program planning with DSW and other Campaign programs to assure requirements needed for certification /qualification are properly set and met.

<sup>7</sup> The ASC Top Ten Risks table was originally published in the *ASC Program Plan FY05*.

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
3	Inability to respond effectively with Modeling & Simulation (M&S) capability and expertise in support of stockpile requirements – near or long term, planned or unplanned (SLEP, SFIs, etc.).	Very High	Low	MEDIUM	Integrate program planning, particularly technical investment priority, with DSW and other Campaign programs to ensure capability and expertise is developed in most appropriate areas; retain ability to apply legacy tools, codes, models.
4	Base of personnel with requisite skills, knowledge, and abilities erodes.	High	Low	MEDIUM	Maintain emphasis on “best and brightest” personnel base, with Institutes, Research Foundations, and University programs, as central feeder elements of the program.
5	Advanced material model development more difficult, takes longer than expected.	Moderate	High	MEDIUM	Increase support to physics research; pursue plans for additional computing capability for physics model development
6	Data not available for input to new physics models or for model validation.	High	Moderate	MEDIUM	Work with Science Campaigns to obtain needed data; propose relevant experiments.
7	Infrastructure resources are insufficient to meet designer, analyst, DSW, or other Campaign program needs.	High	Low	MEDIUM	Integrate program planning with DSW and other Campaigns, to ensure requirements for computing are understood and appropriately set; maintain emphasis on system view of infrastructure and PSE strategy, as central elements of the program.
8	External regulatory requirements delay program deliverables by diverting resources to extensive compliance-related activities	Moderate	Low	MEDIUM	Work with external regulatory bodies to assure that they understand NNSA’s mission, ASC’s mission, and the processes to set and align requirements and deliverables, consistent with applicable regulations.

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
9	Inadequate computational environment impedes development and use of advanced applications on ASC platforms.	Moderate	Very Low	LOW	Integrated planning between program elements to anticipate application requirements and prioritize software tools development and implementation.
10	Fundamental flaws discovered in numerical algorithms used in advanced applications require major changes to application development.	Moderate	Very Low	LOW	Anticipate or resolve algorithm issues through technical interactions on algorithm research through the Institutes, ASC Centers, and academia, and focus on test problem comparisons as part of software development process.



## VIII. Performance Measures

Table VIII-1. ASC Performance Measures

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
<b>Goal:</b> Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
INDICATOR	ANNUAL TARGETS							ENDPOINT TARGET DATE
	FY04	FY05	FY06	FY07	FY08	FY09	FY10	
Peer-reviewed progress in completing milestones, according to a schedule in the ASC Campaign Program Plan, in the development and implementation of improved models and methods into integrated weapon codes and deployment to their users (long-term output). Panel Criteria: (1) Delivery and implementation of validated models into code projects, and (2) Documented verification of approximations.	High Fidelity Primary Code	Initial baseline Primary Code	Initial validated simulation code for W76 and W80	W80 code baseline	Conduct modern baseline of all enduring stockpile systems	Complete modern baseline of all enduring stockpile systems	Quantify margins and uncertainties of modern baseline simulations	By 2015, accomplish full transition from legacy design codes to modern ASC codes with documented quantification of margins and uncertainties of simulation solutions.
Cumulative percentage of the 31 weapon system	32%	38%	51%	67%	87%	96%	100%	By 2010, analyze 100 percent of 31

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
<b>Goal:</b> Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
INDICATOR	ANNUAL TARGETS							ENDPOINT TARGET DATE
	FY04	FY05	FY06	FY07	FY08	FY09	FY10	
components, primary / secondary / engineering system, analyzed using ASC codes, as part of annual assessments and certifications (long-term output).								weapon system components using ASC codes, as part of annual assessments and certifications (interim target).
The maximum individual platform computing capability delivered, measured in trillions of operations per second (teraflops) (long-term output).	40	100	200	200	200	350	350	BY 2009, deliver a maximum individual platform computing capability of 350 teraops.
Total capacity of ASC production platforms attained, measured in teraflops, taking into consideration procurements & retirements of systems (long-term output).	75	172	160	360	470	980	980	By 2009, attain a total production platform capacity of 980 teraops.

<b>ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN</b>								
<b>Goal:</b> Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
<b>INDICATOR</b>	<b>ANNUAL TARGETS</b>							<b>ENDPOINT TARGET DATE</b>
	<b>FY04</b>	<b>FY05</b>	<b>FY06</b>	<b>FY07</b>	<b>FY08</b>	<b>FY09</b>	<b>FY10</b>	
Average cost per teraflops of delivering, operating, and managing all SSP production systems in a given fiscal year (efficiency measure).	\$8.15M	\$5.7M	\$3.99M	\$2.79M	\$1.96M	\$1.37M	\$0.96M	By 2010, attain an average cost of \$0.96 M per teraflops of delivering, operating, and managing all SSP production systems.



## Appendix A. Glossary

---

ASC	Advanced Simulation and Computing
CCF	Central Computing Facility
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
DCE	Distributed Computing Environment
DNT	Defense and Nuclear Technologies Directorate at LLNL
DSW	Directed Stockpile Work
DTA	Design-Through-Analysis
ESN	Enterprise Secure Net
FOUS	Facility Operations and User Support
GPFS	Global Parallel File System
HEDP	High Energy Density Physics
HPC	High Performance Computing
HPSS	High Performance Storage System
I/O	Input/Output
KCP	Kansas City Plant
LANL	Los Alamos National Laboratory
LDAP	Lightweight Directory Access Protocol
LDCC	Laboratory Data Communications Center
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
MPI	Message Passing Interface
NAS	Network-Attached Storage
NFS	Network File System
NIF	National Ignition Facility
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NSA	National Security Agency
nWBS	National Work Breakdown Structure
NWCC	Nuclear Weapons Compute Clusters
PaScalBB	Parallel Scalable Back Bone Concept
PEM	Physics and Engineering Models

PSI	Parallel Storage Interface
QMU	Quantification of Margins and Uncertainties
RAS	Reliability, Availability, and Serviceability
RDMA	Remote Direct Memory Access
RRW	Reliable Replacement Warhead
SCC	Nicolas C. Metropolis Center for Modeling and Simulation
SCN	Secure Classified Network
SFI	Significant Finding Investigation
SLEP	Stockpile Life Extension Program
SNL	Sandia National Laboratories
SRN	Secure Restricted Network
SSP	Stockpile Stewardship Program
STS	Stockpile-to-Target Sequence
TSF	Terascale Simulation Facility
V&V	Verification and Validation
VTK	Visualization Tool Kit
WAN	Wide Area Networking
WBS	Work Breakdown Structure

## Appendix B. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Steve Louis, LLNL, 925-422-1550, <a href="mailto:stlouis@llnl.gov">stlouis@llnl.gov</a> John Thorp, LANL, 505-665-82265, <a href="mailto:thorp@lanl.gov">thorp@lanl.gov</a> Mike Hardwick, SNL, 925-294-2157, <a href="mailto:mfhardw@sandia.gov">mfhardw@sandia.gov</a> Sudip Dosanjh, SNL, 505-845-7018, <a href="mailto:ssdosan@sandia.gov">ssdosan@sandia.gov</a>
1.5.4.1	Capability Systems	Mark Seager, LLNL, 925-423-3141, <a href="mailto:seager1@llnl.gov">seager1@llnl.gov</a> Manual Vigil, LANL, 505-665-1960, <a href="mailto:mbv@lanl.gov">mbv@lanl.gov</a> Sudip Dosanjh, SNL, 505-845-7018, <a href="mailto:ssdosan@sandia.gov">ssdosan@sandia.gov</a>
1.5.4.2	Capacity Systems	Mark Seager, LLNL, 925-423-3141, <a href="mailto:seager1@llnl.gov">seager1@llnl.gov</a> Manual Vigil, LANL, 505-665-1960, <a href="mailto:mbv@lanl.gov">mbv@lanl.gov</a> Sudip Dosanjh, SNL, 505-845-7018, <a href="mailto:ssdosan@sandia.gov">ssdosan@sandia.gov</a>
1.5.4.3	Advanced Systems	Mark Seager, LLNL, 925-423-3141, <a href="mailto:seager1@llnl.gov">seager1@llnl.gov</a> John Thorp, LANL, 505-665-82265, <a href="mailto:thorp@lanl.gov">thorp@lanl.gov</a> Sudip Dosanjh, SNL, 505-845-7018, <a href="mailto:ssdosan@sandia.gov">ssdosan@sandia.gov</a>
1.5.4.4	System Software and Tools	Steve Louis, LLNL, 925-422-1550, <a href="mailto:stlouis@llnl.gov">stlouis@llnl.gov</a> John Thorp, LANL, 505-665-82265, <a href="mailto:thorp@lanl.gov">thorp@lanl.gov</a> Jim Ang, SNL, 505-844-0068, <a href="mailto:jaang@sandia.gov">jaang@sandia.gov</a>
1.5.4.5	I/O, Storage Systems, and Networking	Steve Louis, LLNL, 925-422-1550, <a href="mailto:stlouis@llnl.gov">stlouis@llnl.gov</a> John Thorp, LANL, 505-665-82265, <a href="mailto:thorp@lanl.gov">thorp@lanl.gov</a> Mike Hardwick, SNL, 925-294-2157, <a href="mailto:mfhardw@sandia.gov">mfhardw@sandia.gov</a>
1.5.4.6	Pre- and Post-Processing Environments	Steve Louis, LLNL, 925-422-1550, <a href="mailto:stlouis@llnl.gov">stlouis@llnl.gov</a> John Thorp, LANL, 505-665-82265, <a href="mailto:thorp@lanl.gov">thorp@lanl.gov</a> Dino Pavlakos, SNL, 505-844-9089, <a href="mailto:cjpavla@sandia.gov">cjpavla@sandia.gov</a>
1.5.5	Facility Operations and User Support	Doug East, LLNL, 925-424-4148, <a href="mailto:east1@llnl.gov">east1@llnl.gov</a> Cheryl Wampler, LANL, 505-667-5243, <a href="mailto:clw@lanl.gov">clw@lanl.gov</a> John Zepper, SNL, 505-845-8421, <a href="mailto:jdzepper@sandia.gov">jdzepper@sandia.gov</a>
1.5.5.1	Facilities, Operations, and Communications	Doug East, LLNL, 925-424-4148, <a href="mailto:east1@llnl.gov">east1@llnl.gov</a> Steve Shaw, LANL, 505-606-0203, <a href="mailto:shaw@lanl.gov">shaw@lanl.gov</a> John Noe, SNL, 505-844-5592, <a href="mailto:jpnoe@sandia.gov">jpnoe@sandia.gov</a>
1.5.5.2	User Support Services	Brian Carnes, LLNL, 925-423-9181, <a href="mailto:carnes1@llnl.gov">carnes1@llnl.gov</a> Steve Shaw, LANL, 505-606-0203, <a href="mailto:shaw@lanl.gov">shaw@lanl.gov</a> John Noe, SNL, 505-844-5592, <a href="mailto:jpnoe@sandia.gov">jpnoe@sandia.gov</a>

WBS	Title	Contact
1.5.5.3	Collaborations	Lynn Kissel, LLNL, 925-423-7940, <a href="mailto:kissel1@llnl.gov">kissel1@llnl.gov</a> Cheryl Wampler, LANL, 505-667-5243, <a href="mailto:clw@lanl.gov">clw@lanl.gov</a> Robert Thomas, SNL, 505-844-7450, <a href="mailto:rkthomas@sandia.gov">rkthomas@sandia.gov</a>
1.5.1.4.12	Caltech, Center for Simulating Dynamic Response of Materials	Dan Meiron, 626-395-3424, <a href="mailto:dim@its.caltech.edu">dim@its.caltech.edu</a> Michael Aivazis, 626-395-3424, <a href="mailto:aivazis@caltech.edu">aivazis@caltech.edu</a> Paul Dimotakis, 626-395-6811 Bill Goddard, 626-395-2731, <a href="mailto:wag@wag.caltech.edu">wag@wag.caltech.edu</a> Michael Ortiz, 626-395-4530, <a href="mailto:ortiz@aeor.caltech.edu">ortiz@aeor.caltech.edu</a> Dale Pullin, 626-395-6081, <a href="mailto:dale@galcit.caltech.edu">dale@galcit.caltech.edu</a>
1.5.1.4.12	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, <a href="mailto:moin@stanford.edu">moin@stanford.edu</a> Juan Alonso, 650-723-9954, <a href="mailto:jjalonso@stanford.edu">jjalonso@stanford.edu</a> Heinz Pitsch, 650-736-1995, <a href="mailto:h.pitsch@stanford.edu">h.pitsch@stanford.edu</a> William Dally, 650-725-8945, <a href="mailto:billd@csl.stanford.edu">billd@csl.stanford.edu</a> Pat Hanrahan, 650-723-8530, <a href="mailto:hanrahan@cs.stanford.edu">hanrahan@cs.stanford.edu</a>
1.5.1.4.12	University of Chicago, Center for Astrophysical Flash Phenomena	Donald Q. Lamb, 773-702-7194, <a href="mailto:d-lamb@uchicago.edu">d-lamb@uchicago.edu</a> Robert Rosner, 773-702-0560, <a href="mailto:r-rosner@uchicago.edu">r-rosner@uchicago.edu</a> Todd Dupont, 773-702-3485, <a href="mailto:Dupont@cs.uchicago.edu">Dupont@cs.uchicago.edu</a> Rusty Lusk, 630-252-7852, <a href="mailto:lusk@mcs.anl.gov">lusk@mcs.anl.gov</a> Tomasz Plewa, 773-834-3227, <a href="mailto:tomek@flash.uchicago.edu">tomek@flash.uchicago.edu</a> Andrew Siegel, 773-834-8501, <a href="mailto:siegela@flash.uchicago.edu">siegela@flash.uchicago.edu</a> Jim Truran, 773-702-9584, <a href="mailto:truran@nova.uchicago.edu">truran@nova.uchicago.edu</a>
1.5.1.4.12	University of Illinois, Center for Simulation of Advanced Rockets	Michael T. Heath, 217-333-6268, <a href="mailto:m-heath@uiuc.edu">m-heath@uiuc.edu</a> William A. Dick, 217, 244-7235, <a href="mailto:wdick@uiuc.edu">wdick@uiuc.edu</a> Robert A. Fiedler, 217-333-3247, <a href="mailto:rfiedler@uiuc.edu">rfiedler@uiuc.edu</a>
1.5.1.4.12	University of Utah, Center for Simulation of Accidental Fires and Explosions	David W. Pershing, 801-581-5057, <a href="mailto:David.Pershing@utah.edu">David.Pershing@utah.edu</a> Chuck Wight, 801-581-8796, <a href="mailto:Chuck.Wight@utah.edu">Chuck.Wight@utah.edu</a> Tom Henderson, 801-581-3601, <a href="mailto:tch@cs.utah.edu">tch@cs.utah.edu</a> Philip Smith, 801-585-3129, <a href="mailto:smith@crsim.utah.edu">smith@crsim.utah.edu</a> Patrick McMurtry, 801-581-3889, <a href="mailto:mcmurtry@eng.utah.edu">mcmurtry@eng.utah.edu</a> Eric Eddings, 801-585-3931, <a href="mailto:eddings@che.utah.edu">eddings@che.utah.edu</a> Steve Parker, 801-585-1504, <a href="mailto:sparker@cs.utah.edu">sparker@cs.utah.edu</a>

## **Appendix C.**

### **1.5.1.4.12 Alliance Centers Project (LLNL, LANL, SNL)**

---

The Alliance Centers project, of nWBS product 1.5.1.4 Focused Research, Innovation, and Collaboration, involves various activities in support and management of the ASC Alliance Centers. Primarily, it funds the five Centers at Caltech, Stanford, University of Chicago, University of Illinois at Urbana-Champaign and University of Utah, and technical management support services for the five centers by the tri-lab Alliance Strategy Team (AST) and Technical Support Teams (TSTs). The AST manages and coordinates the project among the labs, five Centers and ASC program office, in terms of planning for and evaluating new solicitations; conducting center reviews and finalizing review panel reports; and publishing documents and organizing workshops or conferences when necessary or appropriate. The TSTs primarily provide support for the centers, in terms of planning for and facilitating interactions between the Centers and the laboratories on an individual researcher basis, visits between the centers and the labs, and yearly spring TST reviews of the centers.

The planned FY07 activities include running a peer-reviewed solicitation for the expected follow-on alliance program (to be called Predictive Science Academic Alliance Program (PSAAP)) starting in FY08, setting up TSTs for the new PSAAP Centers, and conducting the last TST review of each existing center, as well as various other interactions between centers, labs and headquarters as appropriate.

Preliminary planned activities in FY08 include AST and TSTs' interactions with the new PSAAP Centers, organizing a conference whose technical program will include result highlights of the current Alliance Centers program, and other activities as appropriate.

#### *California Institute of Technology, Center for Simulating Dynamic Response of Materials*

The overarching goal of the Caltech ASC Alliance Center is the development of a virtual shock physics facility to simulate dynamic response in materials. An allied goal is to advance the state of the art in modeling using a multiscale paradigm. The research developments of the center are ultimately translated into algorithms that perform scalably on the ASC computing platforms. A significant additional thrust has been the integration of four validation experiments as part of our program. The center's modeling and algorithm development efforts are targeted at the integrated simulation of high velocity impact experiments wherein a detonation elicits shock response in a solid target. The fluid mechanics associated with the detonation simulation is performed using advanced Eulerian advection methods while the solid mechanics of the target are performed using a Lagrangian approach and tetrahedral finite elements. The two approaches are coupled using a novel fluid-solid coupling approach based on level sets. The software environment that supports these simulations comprises the Virtual Test Facility (VTF).

Capabilities include the following:

- Integrated simulation capability: At present the VTF software has the following capabilities: parallel 3D AMR based Eulerian fluid mechanics solver via the AMROC framework, parallel 3D solid mechanics, parallel 3D solid shells, scalable communication and Eulerian-Lagrangian coupling algorithms, multiscale materials models for both solid and fluid. Integrated simulations can be optionally deployed via the Python-based Pyre framework.
- Validation: A key part of the center's research program is the use of four carefully designed and highly diagnosed experiments to validate the VTF software. These experiments are: converging shocks and compressible turbulence experiment – used to validate the shock propagation and turbulence models developed by the center; dynamic deformation – validates the center's multiscale modeling and simulation of high strain rate deformation; dynamic fracture – validates the center's modeling of brittle dynamic fracture; and detonation-fracture – validates the center's integrated simulation capability as well as its modeling of ductile rupture in solids and detonation in combusting gases.
- Research in multiscale modeling: The center sponsors a research program to develop multiscale models for dynamic deformation of metals and high explosives as well as development of models for compressible turbulence in fluids.
- Research in computational science: The center carries out research to develop scalable algorithms to support the integrated simulation. Important contributions here include the fast level set algorithm, a scalable Eulerian-Lagrangian coupling algorithm and the Pyre framework as well as new WENO based advection algorithms for compressible turbulence simulations.

Expected activities and deliverables in FY07 are as follows:

- Integrated simulation capability: We will complete the integration of the AMR solver with the LLNL Paradyn solid mechanics code. We will also complete the integration of multiscale models recently developed in support of the validation experiments described immediately below.
- Validation: The following developments are planned in FY07 for the four validation experiments:
  - Converging shock experiment: Completion of phase 1 experiments with detailed validation simulations using the VTF. Initiation of phase 2 experiments.
  - Dynamic deformation experiment: Characterization of high-strain-rate behavior of tantalum including enhanced full field deformation and temperature measurements; Detailed investigation of the interplay between plastic deformation and phase transformation in polycrystalline  $\alpha$ -iron under shear dominant loading conditions; Development of quantitative error measures for comparison between simulations and experimental results for validation purposes.
  - Detonation fracture experiment: Continue experiments on fluid solid coupling. Carry out detailed measurements of strain and pressure for detonations in tubes with closed ends, bends and tees. Improve detonation generation and measurement methods to reduce uncertainty. New experiments on plastic deformation of tube due to detonations and deflagration to detonation transition.
  - Dynamic fracture experiment: Complete experiments to provide mixed mode and pure mode-II cohesive zone laws of adhesive bonds using asymmetric four-point bending fracture tests. Complete experiments to obtain dynamic cohesive

- zone laws under mode-I and mixed mode loading conditions to verify the sufficiency of quasi-static cohesive zone laws in the VTF simulations
- Multiscale modeling:
    - Initiate investigations of compressible turbulence for fluids governed by more realistic equations of state such as Mie-Gruneisen.
    - Complete a fully-integrated simulation of spallation in aluminum using a multiscale model (from vacancies to void sheets) of porous plasticity (in collaboration with Jarek Knap and Jaime Marian of LLNL).
    - Complete development and integration of porous plasticity model and shear band/localization elements, comparison with observed crack geometries for detonation fracture experiment.
    - Integration of relaxed single crystal plasticity model into polycrystalline calculations and comparison of reconstructed microstructures with experiment  
Extension of phase transition model in Fe to finite temperatures using EoS and elastic moduli data computed from first principles in support of dynamic deformation experiments.
  - Computational science: In FY07 we will continue the development and integration of a graph-based load balancing strategy for AMR to improve scalability. We will also begin development of an implicit version of the AMR solver using a multigrid approach. We also plan to continue to improve pyre. The support for databases will be augmented to support MySQL, which is in wide use among data-intensive scientific computing efforts. Finally, we plan to automate the creation of user interaction screens for configuring components regardless of the mode of operation, from command line driven scripts to graphically driven interfaces, both native and web based.

### *Stanford, ASC Alliance Center for Integrated Turbulence Simulations, CITS*

CITS's mission is the development of a high-fidelity computational framework for the simulation of turbulent thermo-fluid systems involving a variety of physical phenomena and geometrically complex configurations. The simulation environment is based on multiple, integrated codes that address specific physical phenomena, such as detailed turbulence dynamics, multiphase interface tracking, combustion, etc. Strong emphasis is given to the development of efficient and scalable numerical algorithms for the ASC parallel computational platforms.

A complementary goal is the development of streaming supercomputer hardware and supporting software and algorithms; streaming represents the next generation of high performance scientific computing as demonstrated by the recent industrial interest. Recent developments at IBM and the introduction of the CELL processor have spurred renewed interest in this activity. The close similarity between CELL and streaming computers allows one to establish the performance of Streaming computers; the Center's legacy codes are currently being tested on the CELL architecture.

Capabilities include the following:

- **Integrated Simulations.** CITS overarching problem is the simulation of the complete aero-thermodynamic flow path through a jet engine. CITS simulation environment is based on a new version of CHIMPS (Coupler for High-performance Integrated Multi-Physics Simulations), which is more flexible and scalable. This new environment has been extensively verified and is now used for the simulations of the jet engine as well as several other applications. The integrated simulations are built

around the Center's flagship codes, CDP (Large Eddy Simulation approach) and SUmB (based on the Reynolds-averaged Navier-Stokes equations). Considerable efforts have been devoted to the derivation of stable, consistent and accurate multi-code interface boundary conditions. The penalty-based boundary condition treatment, together with symmetric discretization operators allows one to derive stability bounds based on energy estimates.

- **Software Engineering.** Scalability continues to be one of the main drivers to improve the Center's computational tools. Recent full-scale integrated simulations run efficiently on 1000s processors. In addition, a team from CITS has been involved in the unclassified testing of the BGL machine at Livermore. The simulation, one of the largest turbulence simulations ever attempted, contains 35G elements and the results are currently being analyzed to find new answers in the "origin" question (identifying the role of turbulence in planet formation).
- **Physics Modeling.** Accurate turbulence modeling has been the key feature of CITS. In addition models for characterizing combustion, heat release and the pollutant production in the engine combustion chamber have been developed and validated; work on soot modeling is well underway; initial results are promising. Multiphase flow capabilities have also been substantially enhanced by introducing a new paradigm, the Resolved Level Set Grid (RLSG) method, which directly uses the CHIMPS integration software to couple the liquid/gas interface tracking algorithm (based on a Cartesian mesh solver) and CDP. Simulations of liquid jet breakups have been initiated. The RLSG approach naturally allows one to achieve grid convergence for the multiphase interface.
- **Verification and Validation.** CITS has increased the visibility and internal dissemination of Verification and Validation activities to build a more widespread knowledge. Following a 2-day workshop on error estimation and uncertainty quantification, several new activities have been initiated: an adjoint formulation for time periodic flows has been derived to perform error estimation and grid adaption. A new forward-in-time MonteCarlo-based adjoint formulation for purely unsteady flows has also been demonstrated for simple problems. Work on the extension of this approach to more realistic problem is in progress.

Expected deliverables in FY07 include (the first ever) large scale integrated simulation of the entire jet engine flow path from the inlet fan to the exit vanes. The objective is to study the effect of the component interactions on the overall engine performance. Several additional problems are currently on-going to improve the boundary condition treatment (in terms of stability and consistency) at the code interfaces. Adjoint-based approaches for time-periodic and unsteady flows will continue to be a focus of our V&V activities. Applications to realistic configurations will be carried out.

### *University of Chicago, ASC Center for Astrophysical Thermonuclear Flashes*

The goal of the Center is to solve long-standing problems of thermonuclear flashes on the surfaces of compact stars, such as neutron stars (X-ray bursts) and white dwarfs (novae); and particularly, in the interior of white dwarfs (Type Ia supernovae). This remarkable problem includes physical phenomena such as the accretion flow onto the surfaces of these compact stars; shear flow and Rayleigh-Taylor instabilities on the stellar surfaces and interiors; ignition of nuclear burning under conditions leading to convection; and either deflagration or detonation, stellar envelope expansion, and the possible creation of a common envelope binary star system. The Center's scientific goal is realized by means of the construction of a multi-dimensional, multi-physics, simulation code (the "FLASH code"), which is able to carry out numerical simulations of the various aspects of the "FLASH problem." In what follows, we highlight the

capabilities of the code and the important astrophysical studies that have been accomplished using the code. We then specify our expected FY07 deliverables.

The FLASH code is a fully modular, extensible, community code that is capable of simulating a wide variety of problems in astrophysics, laboratory fluid dynamics, and plasma physics. Its capabilities include non-relativistic and relativistic hydro, non-relativistic and relativistic MHD, a variety of equations of state, a variety of nuclear networks, multipole and multigrid self-gravity with both isolated and periodic boundary conditions, massless tracer particles, massive particles for treating dark matter, and diffusive radiation transfer. The code is actively being used by nearly 300 scientists around the world for purposes ranging from algorithm development, education, and hardware testing, to research in computational fluid dynamics and MHD, high energy astrophysics, stars and stellar evolution, and cosmology. The FLASH code has enabled the Center to propose and simulate a self-consistent picture of C/O mixing in the surface layers of white dwarf stars prior to novae outbursts; and to simulate the entire white dwarf star during the deflagration phase of Type Ia supernovae, leading to the discovery of an entirely new and promising mechanism for such supernovae.

The primary scientific activity of the Center in the coming year will be a series of large-scale, integrated, multi-physics simulations of Type Ia supernovae, from ignition through the deflagration and detonation phases, and into the free expansion phase. The primary computational physics and validation activities in the coming year will be the continued development of physics modules in support of these large-scale simulations of Type Ia supernovae, and the initiation of collaborations with the DOE labs to develop new predictive science methodologies. The primary code activities in the coming year will be the continued development of FLASH 3, migration of the FLASH code to new ASC platforms as they become available, and support of the large-scale astrophysics simulations of Type Ia supernovae. The primary visualization activities will be the continued development of FLASHVIEW, and visualization of the large-scale simulations of Type Ia supernovae carried out by the astrophysics group and of validation experiments carried out by the computational physics and validation group. The primary activities of the computer science group will be demonstration of the FLASH code on the next generation of scalable computers. The primary basic physics activities will be the continued study of turbulence and the Rayleigh-Taylor instability for reactive flows.

**Expected Deliverables in FY07 include:**

*Astrophysics:*

- Completion of an improved model for the Type Ia supernova flame that minimizes the effects of curvature;
- A set of 3D simulations of pure deflagration and deflagration plus delayed detonation models of Type Ia supernovae, enabling us to explore the sensitivity of the outcome to initial conditions, including the location and number of ignition points;
- Nuclear reaction network and radiation transfer postprocessing of these simulations, enabling us to determine the nucleosynthetic yields and the light curves.

**Computational Physics and Validation:**

- Simulations of the smoldering phase of novae using the low-Mach number solver that has been developed by the Center;

- Development of predictive science methodologies appropriate to the highly non-linear, multi-physics, multi-scale Type Ia supernova problem
- Uncertainty quantification and sensitivity analyses of the simulations of Type Ia supernovae, using these methodologies and extensive light curve, spectra, and polarization data from observations.

**Code:**

- Release of FLASH 3.0, containing all of the capabilities of FLASH 2 and many new capabilities;
- Migration of the FLASH code to new ASC platforms as they become available; and
- Continued support of the simulations being carried out by astrophysics group.

**Visualization, Computer Science, and Basic Physics:**

- Further development of FLASHVIEW, the FLASH desktop visualization tool;
- Visualization of the large-scale simulations carried out by the FLASH Center;
- Continued development of data handling and transfer software;
- Mathematical and numerical studies of flames and reactive flows; and
- Continued analysis of Eulerian fluid and Lagrangian particle data from the very large simulation of isotropic, homogeneous, driven turbulence that the Center carried out using the LLNL BG/L.

***University of Illinois, Center for Simulation of Advanced Rockets***

The goal of the University of Illinois Center for Simulation of Advanced Rockets (CSAR) is the detailed, whole-system simulation of solid propellant rockets from first principles under both normal and abnormal operating conditions. The design of solid propellant rockets is a sophisticated technological problem requiring expertise in diverse subdisciplines, including the ignition and combustion of composite energetic materials; the solid mechanics of the propellant, case, insulation, and nozzle; the fluid dynamics of the interior flow and exhaust plume; the aging and damage of components; and the analysis of various potential failure modes. Each of these aspects is characterized by very high energy density, extremely diverse length and time scales, complex interfaces, and reactive, turbulent, and multiphase flows.

CSAR is focusing on the reusable solid rocket motor (RSRM) of the NASA Space Transportation System, better known as the Space Shuttle, as its long-term simulation vehicle. The RSRM is a well-established commercial rocket, is globally recognized, and design data and propellant configurations are available. The Center has a Space Act Agreement with NASA in place to share data and simulation results under an Export Control/ITAR relationship. Several smaller scale rockets are also simulated to provide validation data for CSAR codes. Simulations that include full geometric and materials complexity require a sequence of incremental developments—in engineering science, computer science, and systems integration—over an extended period.

Our approach to system integration has been to develop a single executable code containing modules for the various components and an interface code for tying them together. We are following an object-oriented design methodology that hides the data structures and other internal details of the individual component codes. This simplifies development and maintenance of the interface code and the component codes, and also makes it easier to swap different versions of the same component—a critical capability for determining the most efficient algorithms and implementations.

Broadly known as Rocstar, the CSAR simulation code is a fully coupled, multiscale, multiphysics suite of integrated modules for 3D simulation of solid propellant rocket performance on massively parallel computers. The suite is designed to be sufficiently general to solve any fluid/structure interaction problem. Components include fluid dynamics (Rocflo, Rocflu), entrained particle tracking and interaction (Rocpart, Rocsmoke), solid mechanics (Rocsolid, Rocfrac), fracture (Rocfrac), particle packing (Rocpack), combustion (Rocburn, Rocfire), and software interface codes for coupling and mesh association (Rocom, Rocman, Rocface, Rocinteract, and others). Refinement of models reflects the synthesis of fundamental, subscale studies critical for detailed simulations of accident scenarios and for reliable simulation of multiscale phenomena such as combustion and turbulence.

The code has been applied to many research problems in addition to solid propellant rockets. These include validation studies employing super-seismic shocks, human (or animal) arteries, and acoustic interaction between helicopter rotor blades and the helicopter body. Rocpack has recently been licensed to the U.S. Army Engineering Research and Development Center for a preliminary study investigating packing of land mines in rocks and soil, and to a private company studying emulsion technologies.

CSAR-NNSA/DP Interactions: Center personnel have traveled extensively and have been involved in a large number of technical and informational meetings. These included meetings intended to explore rocket science and technology, identify technical collaborators, describe the ASC/ASAP program, and establish relationships among Center investigators, DOE lab scientists, and industry leaders. One of the leading topics of discussion between lab staff and CSAR investigators has been and will continue to be in pursuing the potential use within the labs of specific technologies developed at CSAR. Immediate technology candidates include our technology for data transfer at component interfaces and our framework for integration of separately developed codes with automated load balancing.

Student-NNSA/DP Interactions: To date, 33 former CSAR students and five staff members have joined the DOE/NNSA labs as permanent employees (in FY06, nine students and one staff member were added to the list). CSAR has been remarkably successful in encouraging student-lab interactions. Opportunities for UIUC graduate and undergraduate student interaction with the NNSA/DP laboratories include summer student internships, joint research, and CSAR undergrads in labs collaborating in research with NNSA/DP scientists.

Expected Deliverables in FY07:

- RSRM simulations — Semi-annual simulations of RSRM exercise the then-most recent features of Rocstar. RSRM simulations will increase in complexity and machine demand through the end of DOE/NNSA support, concluding with a 10,000+/- processor simulation runs in 2007.
- RSRM complete normal burn — The ignition transients for the Space Shuttle booster are well characterized in the open literature, and we have access to extensive test data. An especially difficult aspect of simulating the entire history of a large motor is reducing the run time. For a fluids mesh that is fine enough to allow accurate turbulence modeling, for example, time zooming techniques under consideration will be required to reach 120 seconds of physical problem time. We have two distinct methodologies for time zooming; both will be exercised in FY06 and 07.
- Titan IV case rupture accident — In this simulation, the pressure builds up until the case fails. In test firing an early design (1991), a rocket motor exploded violently destroying the test stand, but there was no propellant detonation. This simulation

will include the use of the new advanced material model for the propellant that includes the effects of voids and dewetting.

- Other major rocket simulations include aluminized propellant combustion and burnout in the U.S. Air Force ballistic test system (BATES) motor; an end-to-end validation study using a commercial attitude control motor (Aerojet); smoke, temperature and vorticity validation; and an enhanced simulation of the Titan IV case rupture accident.
- Other efforts include:
  - Study acoustic effects of 3D vortex shedding in RSRM simulations downstream of intersegment inhibitors
  - Validate algorithm for propellant regression along case
  - Implement material model and new finite element for case insulation
  - Implement implicit code and test in Rocflu module
  - Develop large-scale motor (RSRM and Titan) meshes for use on machines with 10,000 or more processors

#### *University of Utah, Center for Simulation of Accidental Fires and Explosions*

The principal objectives of C-SAFE are to advance the state of the art in high-performance computer simulation science by creating world-class software and by educating students, postdocs and professional staff in the science and art of high-performance simulation. C-SAFE software is particularly suited for creating, verifying, validating and visualizing simulations of complex physical and chemical behavior over wide ranges of time and space. The C-SAFE target scenario is to simulate the explosive response of a cylindrical steel container of PBX9501 embedded in a 10-m diameter jet fuel fire. A key goal is to create multi-physics simulation software for which the behavior at the resolved (grid) scale accurately reflects underlying physics associated with smaller unresolved scales, down to the scales of individual atoms and molecules.

The C-SAFE target simulation scenario places stringent demands on the software in order to produce validated results. The fire simulation uses Large Eddy Simulation (LES) on a structured grid to capture the effects of radiant and convective heating on the length and time scales relevant to large pool fires such as may be present after an airplane crash or other transportation accident. The Center has developed the methodology for identifying chemical surrogates for jet fuel and other complex hydrocarbon mixtures, as well as detailed kinetic mechanisms for combustion in air. Because energy transport in fuel fires is dominated by radiation from soot particles, C-SAFE has created a world-class center for understanding the chemistry of soot formation and has pioneered techniques for incorporating that knowledge into large-scale simulations.

Although the fire simulation occurs on time scales of seconds to minutes, the resulting explosion occurs over the course of a few milliseconds. The code therefore incorporates methods of detecting the rare trigger events and executing appropriate changes in time stepping, while incorporating only those physics models that are relevant to the process and time scale of interest. Creating validated simulations of explosions has required the development of new mechanics codes (such as, an extension of the Material Point Method) for simulating closely coupled fluid-structure interactions on multiple time scales in ways that allow seamless transitions between compression, metal plasticity, fracture and explosive release of gases at arbitrary locations throughout the simulation domain.

The chemical and mechanical properties of plastic-bonded explosives have been investigated from first principles using molecular dynamics simulations, the results of which are used to formulate validated homogeneous computational models and constitutive properties for use at the resolved grid scale of the simulations.

All of these capabilities are combined in a single scalable software infrastructure (Uintah Core Code Development), which allows large-scale interactive visualization, incorporation of adaptive mesh refinement (AMR) algorithms and infrastructure, as well as continuous code verification through daily testing. The code has been ported to five major supercomputer architectures and exhibits excellent scaling characteristics to at least 2000 processors.

The principal scientific application of the code is analysis of heating rates and fire geometries on the violence of the resulting explosion (such as, limiting behaviors of fast and slow cookoff, as well as intermediate behaviors and transitions). Because the computational modules are designed to reflect underlying fundamental physics rather than empirical observations, the code also has proved capable of simulating many different types of complex processes such as cell membrane mechanics and interaction of projectiles with human tissue. Therefore, the intrinsic capabilities that were built into the software for one purpose can be leveraged for many different types of scientific investigations driven by high-performance computer simulations.

Expected deliverables in FY07 include:

- Full implementation of repeatable and reversible changes in simulation algorithms to handle physics on disparate timescales
- Implementation of a new computational model for unsteady combustion of energetic materials
- Validation simulations of heating of large containers in rectangular pool fires
- Semi Implicit ICE-AMR JP-8 pool fire simulation
- Semi-implicit MPMICE-AMR calculation involving reacting flow
- End-to End simulation using multi-resolution computational grid



