Advanced Simulation & Computing™

October 23, 2008

NA-ASC-116R-08-Vol. 2, Rev. 0.5
LLNL-TR-407766

FY09–FY10 Implementation Plan

Volume 2, Rev. 0.5

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

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

Advanced Simulation and Computing

FY09–10 IMPLEMENTATION PLAN
Volume 2, Rev. 0.5

October 23, 2008

Approved by:

Robert Meisner, NNSA ASC Program Director (acting)	Signature	*10/23/08* Date
James Peery, SNL ASC Executive	Signature	*10/13/08* Date
Michel McCoy, LLNL ASC Executive	Signature	*10/20/08* Date
John Hopson, LANL ASC Executive	Signature	*10/20/08* Date

ASC Focal Point
Robert Meisner
NA 121.2
Tele.: 202-586-1800
FAX: 202-586-0405
bob.meisner @nnsa.doe.gov

IP Focal Point
Njema Frazier
NA 121.2
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 Roadmap Drivers for FY09–FY10			
VII.	ASC Risk Management			
VIII.	Performance Measures			
IX.	Budget			
Appendix A.	Glossary			
Appendix B.	Codes			
Appendix C.	Points of Contact			
Appendix D.	1.5.1.4-TRI-001 Academic Alliance Centers			
Appendix E.	ASC Obligation/Cost Plan			
Appendix F.	ASC Performance Measurement Data for FY09			

Contents

I. EXECUTIVE SUMMARY 1

II. INTRODUCTION 2

ASC Contributions to the Stockpile Stewardship Program 3

III. ACCOMPLISHMENTS FOR FY07–FY08 7

Computational Systems and Software Environment 7
Facility Operations and User Support 9
Academic Alliances 10

IV. PRODUCT DESCRIPTIONS BY THE NATIONAL WORK BREAKDOWN STRUCTURE 13

WBS 1.5.4: Computational Systems and Software Environment 13

WBS 1.5.4.1: Capability Systems 13
WBS 1.5.4.1-LLNL-001 Purple 14
WBS 1.5.4.1-LLNL-002 Sequoia 14
WBS 1.5.4.1-LANL-001 Systems Requirements and Planning 15
WBS 1.5.4.1-SNL-001 Red Storm Capability Computing Platform 15
WBS 1.5.4.1-LANL/SNL-001 Alliance for Computing at Extreme Scale Zia Capability Computing Platform 16
WBS 1.5.4.1-LANL/SNL-002 Alliance for Computing at Extreme Scale Architecture Office 17
WBS 1.5.4.2: Capacity Systems 17
WBS 1.5.4.2-LLNL-001 Tri-Lab Linux Capacity Cluster 18
WBS 1.5.4.2-LANL-001 Capacity System Integration 19
WBS 1.5.4.2-SNL-001 ASC Capacity Systems 19
WBS 1.5.4.3: Advanced Systems 20
WBS 1.5.4.3-LLNL-001 BlueGene/P and BlueGene/Q Research and Development 20
WBS 1.5.4.3-LLNL-002 Petascale Application Enablement 21
WBS 1.5.4.3-LANL-002 Roadrunner Phase 3 Delivery and Acceptance 22
WBS 1.5.4.3-SNL-001 Advanced Systems Technology Research and Development 22
WBS 1.5.4.4: System Software and Tools 23
WBS 1.5.4.4-LLNL-001 System Software Environment for Scalable Systems 24
WBS 1.5.4.4-LLNL-002 Applications Development Environment and Performance Team 25
WBS 1.5.4.4-LANL-001 Roadrunner Computer Science 26
WBS 1.5.4.4-LANL-002 Software Technologies for Next Generation Platforms 26
WBS 1.5.4.4-LANL-003 Code Performance and Throughput 28
WBS 1.5.4.4-LANL-005 Software Support 29
WBS 1.5.4.4-LANL-006 Applications Readiness 30
WBS 1.5.4.4-LANL-007 Productivity Project 30
WBS 1.5.4.4-SNL-001 Software and Tools for Scalability and Reliability Performance 31
WBS 1.5.4.4-SNL-003 System Simulation and Computer Science 32
WBS 1.5.4.5: Input/Output, Storage Systems, and Networking 33

WBS 1.5.4.5-LLNL-001 Archive Storage....34
WBS 1.5.4.5-LLNL-002 Parallel and Network File Systems....35
WBS 1.5.4.5-LLNL-003 Networking and Test Beds....36
WBS 1.5.4.5-LANL-001 File Systems and Input/Output Project....36
WBS 1.5.4.5-LANL-002 Archival Storage Design and Development....37
WBS 1.5.4.5-SNL-001 Scalable Input/Output and Storage Systems....38
WBS 1.5.4.5-SNL-003 Archival Storage....39
WBS 1.5.4.5-SNL-004 Scalable Input/Output Research....39
WBS 1.5.4.6: Post-Processing Environments....40
WBS 1.5.4.6-LLNL-001 Scientific Visualization....41
WBS 1.5.4.6-LLNL-002 Scientific Data Management....42
WBS 1.5.4.6-LANL-001 Visualization and Insight for Petascale Simulations Project....43
WBS 1.5.4.6-LANL-002 Production Systems for Visualization and Insight Project....44
WBS 1.5.4.6-LANL-003 Physics-Based Simulation Analysis Project....45
WBS 1.5.4.6-SNL-001 Remote Petascale Data Analysis....46
WBS 1.5.4.6-SNL-002 Visualization Deployment and Support....47
WBS 1.5.4.7: Common Computing Environment....48
WBS 1.5.4.7-TRI-001 Tripod Operating System Software....49
WBS 1.5.4.7-TRI-002 Open|SpeedShop....49
WBS 1.5.4.7-TRI-003 Workload Characterization....50
WBS 1.5.4.7-TRI-004 Application Monitoring....51
WBS 1.5.4.7-TRI-005 Shared Work Space....51
WBS 1.5.4.7-TRI-006 Gazebo Test and Analysis Suite....52
WBS 1.5.5: Facility Operations and User Support....55
WBS 1.5.5.1: Facilities, Operations, and Communications....55
WBS 1.5.5.1-LLNL-001 System Administration and Operations....56
WBS 1.5.5.1-LLNL-002 Software and Hardware Maintenance, Licenses, and Contracts....56
WBS 1.5.5.1-LLNL-003 Computing Environment Security and Infrastructure....57
WBS 1.5.5.1-LLNL-004 Facilities Infrastructure and Power....58
WBS 1.5.5.1-LLNL-005 Classified and Unclassified Facility Networks....59
WBS 1.5.5.1-LLNL-006 Wide-Area Classified Networks....60
WBS 1.5.5.1-LANL-001 High-Performance Computing Operations Requirements Planning....60
WBS 1.5.5.1-LANL-002 Roadrunner Phase 3 Initial Deployment....61
WBS 1.5.5.1-LANL-003 Ongoing Network Operations....61
WBS 1.5.5.1-LANL-004 Network Infrastructure Integration....62
WBS 1.5.5.1-LANL-005 Ongoing Systems Operations....63
WBS 1.5.5.1-LANL-008 Ongoing Facilities....64
WBS 1.5.5.1-SNL-001 Production Computing Services....65
WBS 1.5.5.1-SNL-002 Facilities and Infrastructure....66
WBS 1.5.5.1-SNL-003 Tri-Lab System Integration and Support....66
WBS 1.5.5.1-Y12-001 Applications in Support of Manufacturing Production and Connectivity....67
WBS 1.5.5.2: User Support Services....68
WBS 1.5.5.2-LLNL-001 Hotlines and System Support....69
WBS 1.5.5.2-LANL-001 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support....69
WBS 1.5.5.2-SNL-001 User Environment and Application Support....70
WBS 1.5.5.3: Collaborations....71
WBS 1.5.5.3-LLNL-001 Program Support....72
WBS 1.5.5.3-LLNL-002 Scientific Collaborations....72
WBS 1.5.5.3-LANL-001 Program Support....73
WBS 1.5.5.3-SNL-001 One Program/Three Labs....73

V. ASC LEVEL 1 AND 2 MILESTONES 75

VI. ASC ROADMAP DRIVERS FOR FY09–FY10 97

VII. ASC RISK MANAGEMENT 99

VIII. PERFORMANCE MEASURES 103

APPENDIX A. GLOSSARY 105

APPENDIX C. POINTS OF CONTACT 109

APPENDIX D. WBS 1.5.1.4-TRI-001 ACADEMIC ALLIANCE CENTERS 113

University of Chicago 114
University of Illinois at Urbana-Champaign 115
University of Utah 116
California Institute of Technology 117
Purdue 119
Stanford University 120
University of Michigan 121
University of Texas 122

APPENDIX E. ASC OBLIGATION/COST PLAN 125

APPENDIX F. ASC PERFORMANCE MEASUREMENT DATA FOR FY09 127

List of Tables

Table V-1. Quick Look: *Proposed* Level 1 Milestone Dependencies....75
Table V-2. Quick Look: Level 2 Milestone Dependencies for FY09 76
Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY10....78
Table VI-1. ASC Roadmap Drivers for FY09-10 97
Table VII-1. ASC's Top Ten Risks....99
Table VIII-1. ASC Campaign Annual Performance Results and Targets 103
Table E-1. ASC Performance Measurement Data for FY09.... 127

List of Figures

Figure D-1. ASC obligation/cost plan for FY08. 125

I. Executive Summary

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the surety and reliability of the U.S. nuclear stockpile. The SSP uses past nuclear test data along with current and future non-nuclear 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 requires 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)[1] 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 (LEPs) 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 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 performance 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.

[1] 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) overarching goal of Nuclear Weapons Stewardship: *"We continue to advance the Stockpile Stewardship Program to push the scientific and engineering boundaries needed to maintain our nuclear arsenal. It also means maintaining the basic science and engineering that is the foundation of the weapons program."*[2]

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 Thomas P. D'Agostino summarized[3] the NNSA objectives for SSP as follows:

"Our fundamental national security responsibilities for the United States include:

- *Assuring the safety, security and reliability of the U.S. nuclear weapons stockpile while at the same time transforming the stockpile and the infrastructure that supports it;*
- *Reducing the threat posed by nuclear proliferation; and,*
- *Providing reliable and safe nuclear reactor propulsion systems for the U.S. Navy."*

"Throughout the past decade, the Stockpile Stewardship Program (SSP) has proven its ability to successfully sustain the safety, security and reliability of the nuclear arsenal without resorting to underground nuclear testing. The SSP also enables the U.S. to provide a credible strategic deterrent capability with a stockpile that is significantly smaller. To assure our ability to maintain essential military capabilities over the long-term, however, and to enable significant reductions in reserve warheads, we must make progress towards a truly responsive nuclear weapons infrastructure as called for in the Nuclear Posture Review (NPR). The NPR called for a transition from a threat-based nuclear deterrent, with large numbers of deployed and reserve weapons, to a deterrent that is based on capabilities, with a smaller nuclear weapons stockpile and

[2] NNSA Strategic Planning Guidance for FY2010–2014, April 2008, page 17.

[3] Testimony on the *FY 2008 National Defense Authorization Budget Request for the Department of Energy's NNSA* before the House Armed Services Subcommittee, March 20, 2007.

greater reliance on the capability and responsiveness of the Department of Defense (DoD) and NNSA infrastructure to adapt to emerging threats."

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;[4] to restore production 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) 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 enable QMU and to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW provides requirements for simulation, including planned LEPs, 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 (HPC) 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-teraFLOPS systems in the world and have both since been decommissioned.

[4] While there are no plans to develop new weapons, acquiring such capability is an important prerequisite to deep reductions in the nuclear stockpile.

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-teraFLOPS 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-teraFLOPS platform on the ASCI Q machine, which has been retired from service.

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 non-nuclear experiments and diagnostics.

In FY05, ASC identified and documented SSP requirements to move beyond a 100-teraFLOPS 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.

In FY06, ASC delivered 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 were demonstrated through 2D and 3D physics and engineering simulations. The engineering simulations analyzed system behavior in abnormal thermal environments and mechanical response of systems to hostile blasts. Additionally, confidence measures and methods for uncertainty quantification (UQ) were developed to support weapons certification and QMU Level 1 milestones.

In FY07, ASC supported the completion of the W76-1 and W88 warhead certification, using quantified design margins and uncertainties; ASC also provided two robust 100-teraFLOPS-platform production environments by IBM and CRAY, supporting DSW and Campaign simulation requirements, respectively. One of the original ASCI program Level 1 milestones was completed when the ASC Purple system was formally declared "generally available." This was augmented by the 360-teraFLOPS ASC BlueGene/L system, which provided additional capability for science campaigns. The ASC-funded partnerships between SNL/Cray and LLNL/IBM have transformed the supercomputer industry. There are currently at least 34 "Blue Gene Solution" systems on the Top 500 list and 38 Cray sales based on the SNL Red Storm architecture.

In FY08, ASC delivered 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, named Roadrunner, was sited at LANL. SNL and LANL established the collaborative Alliance for Computing at Extreme Scale (ACES) for the purpose of providing a user facility for production capability computing to the Complex. Plans were made for machine Zia, the first machine to be hosted through ACES, to be procured and sited at LANL.

By FY09, a suite of physics-based models and high-fidelity databases will be developed and implemented. ASC is being brought to bear on critical simulations in support of secure transportation and NWC infrastructure.

In FY10 and beyond, ASC will continue to deliver codes to address the next generation of LEPs and 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. In addition, ASC will continue to provide national leadership in HPC and deploy capability and capacity platforms in support of Defense Programs campaigns.

III. Accomplishments for FY07–FY08

ASC accomplishments from Quarter 4, fiscal year 2007, through quarter 3, fiscal year 2008, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Received signature and approval for Sequoia Critical Decision 1 (CD1) by NNSA Deputy Administrator Robert Smolen on March 26, 2008. CD1 is the Conceptual Baseline package for the project (covering scope, cost, schedule, requirements, impact, and acquisition strategy) and is the follow-on to the CD0 Mission Need package approved in April 2007.
- Successfully expanded size of BlueGene/L by about 40 percent (to 596 teraFLOPS peak from 367 teraFLOPS, to 70 tebibytes of memory from 32 tebibytes, to 212,992 processing cores from 131,072), keeping BlueGene/L as number one on the November 2007 Top 500 list.
- Completed the Level 2 milestone to deploy Moab resource management services on BlueGene/L. This milestone allows batch jobs on BlueGene/L to now be submitted, scheduled, and run through the Moab Workload Manager (the standard tri-lab batch scheduling system).
- Implemented Tripod operating system software (TOSS) governance structure and the support infrastructure, including a software repository—Bugzilla—and documentation. Synthetic workload testing of Tri-Lab Linux Capacity Cluster (TLCC) test systems was successfully completed at the vendor, and TLCC systems were delivered to all three labs. TOSS is now deployed on TLCC platforms.

LANL Accomplishments for Computational Systems and Software Environment

- A final system technical assessment of Roadrunner was conducted in October 2007. After a successful review, the option to purchase Phase 3 of the Roadrunner system was exercised. This met a Level 2 milestone for FY08.
- Procured a hybrid advanced architecture for the final Roadrunner system.
- The Roadrunner Phase 3 system exceeded a sustained petaFLOPS while running an industry benchmark and became the June 2008 Top 500 #1 computing platform.
- Created the NNSA ACES between LANL and SNL, devoted to providing high-performance capability computing assets required by NNSA's stockpile stewardship mission.
- In an important collaboration with the other labs, LANL completed the *Petascale Infrastructure Plan* milestone.
- Deployed Tripod stack and tools on TLCC systems at LANL.

SNL Accomplishments for Computational Systems and Software Environment

- Developed a new version of Red Storm's Catamount operating system, Catamount N-Way, in support of the FY08 Q4 upgrade. In addition to providing support for the AMD quad-core Opteron, Catamount N-Way also incorporates a power-saving feature to reduce processor power consumption during idle states and an improved shared memory transport, SMARTMAP, for reduced intra-node message passing interface (MPI) latencies.
- During the fourth quarter of FY08, the Red Storm system was outfitted with increased computational capacity from approximately 124 teraFLOPS to approximately 284 teraFLOPS and increased storage capacity from approximately 340 terabytes to 1.9 petabytes. This upgrade allows Red Storm to continue to provide a production-quality capability computational resource for ASC applications and potential move into SCI applications.
- Increased Red Storm's disk capacity by 1.5 petabytes, bringing total capacity to 1.84 petabytes. Applications can now produce higher resolution output, consuming all of the compute node memory, and have sufficient disk capacity for result and checkpoint/restart files. This translates into increased application productivity.
- Hosted the inaugural Institute for Advanced Architectures and Algorithms (IAA) workshop, Memory Opportunities for HPC. The workshop addressed issues associated with the ever-increasing gap between compute and memory performance in processor architectures. Participants were from the government agencies, national labs, industry and academia.
- Demonstrated parallel Structural Simulation Toolkit (SST) prototype. This demonstration examined many of the issues in creating a full parallel version, which will become the foundation for the IAA simulation strategy.
- Developed data analysis features for ParaView. The new features provide scalable analysis capabilities required to support verification and validation (V&V) of large data, with particular emphasis on the comparison of simulation and test data. These new capabilities were used to identify and characterize fragments in extremely large simulation results from CTH, a shock physics code. This capability enables the results of a CTH simulation to be compared with other data, and it allows for simulation results to be propagated to other codes.
- Analyzed and visualized results from a national security calculation that ran on the entire Red Storm platform (more than 25,000 processors) for two months using specific feature-characterization capabilities introduced into ParaView.
- Developed new operational procedures to address and repair file-system consistency issues related to hardware redundant array of independent disks on Red Storm. These repairs previously required large amounts of hands-on consultation to fix. We also designed and implemented a read-only mount capability within the Catamount-SYSIO library to prevent a total loss of access while addressing file-system consistency issues.
- Completed the tri-lab Level 2 milestone to produce the Petascale Environments document on time. This document is helping guide the efforts and priorities in achieving an integrated petascale environment for all of the tri-lab sites.
- Collaborated with industry partners to research and evaluate the impact of adaptive routing and congestion control in large switching fabrics on MPI and input/output (I/O) performance. This technology may be the only hope to create effective large-scale switches for the global file systems serving the ACES Zia and Sequoia

environments. The Woven technology we helped improve won HPCwire Editors' Choice Award for "Most important HPC networking product or technology" at the SC07 conference.

Tri-Lab Accomplishments for Computational Systems and Software Environment

- Completed the Level 2 milestone Infrastructure Plan for ASC Petascale Environments. This formal planning document identifies, assesses, and specifies the development and deployment approaches for ASC petascale platform infrastructures.
- Released TOSS version 1 on schedule, March 3, 2008, for general availability (GA). All three labs passed the synthetic workload acceptance tests for their respective TLCC systems using the TOSS software stack.
- Developed requirements and prototypes for application monitoring tool.
- Defined tri-lab requirements for the shared workspace environment.
- Used the Gazebo acceptance test package as the underlying infrastructure for testing each of the new TLCC systems. During the TLCC acceptance phase, the test package was used to quantify system node utilization and test coverage, leading to improved system stabilization.

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- Rolled out TOSS v1 on our existing Linux cluster installed base in addition to TLCC. This is important as it continues to reduce the total cost of ownership associated with maintaining existing systems. Security patches, bug fixes, and Lustre software integration were accomplished on only one version of the software stack for the entire installed base of capacity Linux clusters.
- Completed Terascale facility 3-MW upgrade. This included 1.5 MW of 480V, an industry power efficiency best practice. This upgrade positioned the center to provide power for the Dawn initial delivery (ID) system as well as out-year capacity clusters.
- Designed and deployed an ESNet backup connection between LANL and LLNL using the DISCOM wide area network (WAN). This connection was actually used one week after it was deployed, when a fiber cut severed LANL's normal ESNet Internet connection for a day and the traffic was rerouted to California. Should LLNL ever have an outage, our traffic can be routed to New Mexico. This development provides a reliable redundant path and saved the cost of paying for an alternate ISP connection.
- Completed design of new identity management system. The identity management system is a critical cost saving measure as it will automate the creation, maintenance, and deletion of user accounts at the center and eliminate processing of paper forms associated with those functions. The design includes electronic workflow of business logic including account approval mechanisms as well as integration with institutional people databases.

LANL Accomplishments for Facility Operations and User Support

- Upgraded equipment in the Nicolas C. Metropolis Center for Modeling and Simulation (SCC) to increase the electrical power SCC computer room to 9.6 MW. This upgrade was in preparation for the Roadrunner Phase 3 system.

- Decommissioned and removed the QB, CA, CB, CC, and QSC clusters (sections of the ASC Q system).
- Received and integrated the first TLCC cluster, Lobo, into the LANL Turquoise network. Lobo is running the Tripod TOSS software stack.
- Designed, developed, and deployed the initial HPC monitoring infrastructure on Roadrunner Phase 1 and the LANL TLCC platforms

SNL Accomplishments for Facility Operations and User Support

- Upgraded classified Red Storm to Red Rose connectivity to the full fifty 10GEthernet paths using the largest production 10GEthernet switches in industry. The increased bandwidth of 30 GB/s improves data transfer to LLNL and LANL.
- Upgraded Lustre routers serving all of the unclassified capacity computing resources from 32 1GEthernets to 4 10GEthernets and 64 bit nodes. Increased available storage by 1.4 petabytes.
- Placed classified and unclassified StorageTek SL8500 tape libraries into production, which is an initial step in reducing from four to two high-performance storage system (HPSS) systems.
- Deployed TLCC systems for testing and early users.
- Outfitted the Red Storm system with increased computational capacity rising from approximately 124 teraFLOPS peak to approximately 284 teraFLOPS peak theoretical speed. The installed memory was upgraded to a uniform 2 gigabytes per processor during this upgrade.

Academic Alliances

University of Chicago Accomplishments

- Released the final version of FLASH 3.0, a highly capable, fully modular, extensible, community code. The FLASH code has now been downloaded more than 1700 times and has been used in more than 320 published papers.
- Published results on 3D isotropic, homogeneous, weakly compressible, driven turbulence. The Flash Center simulation, which was run on the BlueGene/L, is the first one large enough to discriminate among models of turbulence.
- Completed 3D simulations, analysis, and publication of shock-driven R-T and R-M validation experiments.
- Initiated extensive verification studies of buoyancy-driven turbulent nuclear burning.
- Extended its large-scale, 3D simulations of the "gravitationally confined detonation" mechanism for SNe Ia to new initial conditions and through the detonation phase of the explosion.

University of Illinois at Urbana-Champaign Accomplishments

- Completed full burn-out simulation of Reusable Solid Rocket Motor (RSRM)—the long-term goal simulation for CSAR
- Implemented and validated "time zooming" for RSRM and smaller solid rocket motors, accelerating the time-to-completion for solid rocket motor simulations by up to 50 times

- Extended and validated non-uniform-shaped particle packing simulations using micro-tomographic imaging of commercial solid propellants
- Published seminal studies of propellant burning studies that describe how propellant morphology affects amplitude and frequency of vortex shedding
- Aided NASA and U.S. rocket companies in assessing the magnitude and frequency of acoustic vibrations in Ares Constellation (RSRMV) solid rocket launch vehicle

University of Utah Accomplishments

- Completed 16 end-to-end simulations using the Uintah software. This set of simulations predicts the response of an energetic device (time-to-explosion and explosion violence) to heating by jet fuel fires.
- Demonstrated the scalability of Uintah on large adaptive calculations involving AMR with very light computational tasks, which showed that the algorithms and the code scale in an automated way and have low overhead up to 8K processors.
- Validated the Implicit MPM formulation. Comparisons were made with analytical solutions and finite element calculations for heat conduction in composite materials that highlighted a subtle error in boundary conditions at the interface between materials, which was subsequently resolved.
- Verified and validated simulation results to include the quantification of the verification error as applied to validation data from both helium plumes from SNL, and to JP-8 fires from SNL, the University of Nevada at Reno and ATK.
- Implemented the Hugoniostat Simulation Methodology, which allows for efficient atomistic level simulation studies of the energetics and deformation structures, including yielding, of shocked materials.

Stanford University Accomplishments

- Completed the full-system multi-code simulations for the entire flow/thermal path of a Pratt & Whitney 6000 engine. Close collaboration with Pratt & Whitney engineers has led to the transfer of the entire computational dataset (2 Tb) for in-depth analysis to P&W and United Technologies Research Center.
- Transferred CHIMPS, Stanford's integration tool and key ingredient of several new research activities, to NASA and other commercial entities.
- Developed a novel manufactured solution for variable-density, reacting flows.
- Continued the enhancement of the SBP/SAT formulation, a discretization that mimics the symmetry properties of the continuous differential operators in the governing equations, to extend to fully compressible equations on unstructured grids, thus creating an entire suite of tools that embrace this methodology across flow regimes. This effort will become the starting point of the computational infrastructure for the new Predictive Science Academic Alliance Program (PSAAP) Center.
- Developed Sequoia, a programming paradigm with the objective of facilitating efficient scientific computations on conventional clusters as well as the new multi- and many-core processors, and implemented the Center's code on novel architectures, including graphical processing units (GPUs).

California Institute of Technology Accomplishments

- Completed the Phase 0 experiments in which we simulated Mach reflection of a shock as it enters a converging geometry. We have done this for two gases: Nitrogen and CO2. The Virtual Test Facility (VTF) remains predictive upon reshock.
- Characterized the high-strain-rate behavior of tantalum, including enhanced full-field deformation and temperature measurements.
- Performed experiments on impulse wave loading in steel tubes in support of work under an Office of Naval Research (ONR) Multidisciplinary University Research Initiative (MURI).
- Made significant progress as regards to the development of fluid solvers for more complex and realistic equations of state such as Mie-Gruneisen.
- Completed a fully integrated simulation of spallation in aluminum using a multiscale model (from vacancies to void sheets) of porous plasticity.
- Completed development and integration of porous plasticity model and shear band/localization elements.
- Developed an implicit version of the AMR solver using a multigrid approach. This is a very important development as we can now more efficiently simulate turbulence in compressible flows.

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 I/O, storage and networking, to post-processing visualization and data analysis tools, and to a common computing environment. 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 FY09

- Sequoia Performance Baseline/Construction Readiness CD2/3 package
- Sequoia ID (Dawn) delivery, installation, and initial science runs
- Roadrunner Phase 3 final system delivery, acceptance testing, and initial deployment including network integration
- Award contract for the Zia system
- Partial delivery of the Zia system by September 2009
- Quantified actual throughput gain of an NNSA application using the upgraded Red Storm quad-core processors

WBS 1.5.4.1-LLNL-001 Purple

The Purple project previously delivered the IBM Purple cluster of pSeries POWER5-based symmetric multiprocessors at LLNL. The Purple system was declared GA in December 2006 and is currently fully subscribed delivering Capability Computing Campaign (CCC) capabilities to the ASC customers. Expected system lifetime is five years. This project is currently in production operational mode and receives funding for ongoing IBM maintenance.

Planned activities in FY09:

- Continue Purple system maintenance as required

Expected deliverables in FY09:

- Production capability cycles as per CCC guidance

Preliminary planned activities in FY10:

- Continue Purple system maintenance as required

WBS 1.5.4.1-LLNL-002 Sequoia

The Sequoia project will deploy a multi-petaFLOPS computer in late 2011 or early 2012 to be operated as an SSP user facility focused on supporting UQ and reduction in phenomenology (the elimination of code "knobs"). The primary missions of the machine will be (1) UQ for certification and model validation; and (2) weapons science investigations whose resolution is necessary for predictive simulation and, therefore, stockpile transformation. Sequoia will provide computational resources up to 24 times more capable than ASC Purple for UQ and up to 50 times more capable than BlueGene/L for weapons science investigations. Sequoia will bridge the gap between current terascale systems and later exascale systems that will become available within a decade. The Sequoia acquisition strategy also requests technology R&D roadmaps from vendors, with intent that the project will fund an R&D effort for the winning vendor to address identified risks and issues for the platform build and delivery.

There are two major deliverables for Sequoia. The first deliverable (acquisition and delivery of an early environment—called Dawn) is to be completed in 2009, and the second deliverable (acquisition and delivery of final Sequoia environment) is to be completed by end of calendar 2011 or early 2012. These environments (both Dawn and Sequoia) will consist of a large compute platform, plus requisite federated switch networking infrastructure and parallel file system storage hardware (augmenting LLNL's existing Lustre parallel file system deployments) to support compute platforms. Acquired switching infrastructure and storage hardware may also have high-speed hardware connectivity to servers and resources at LLNL outside of Dawn and Sequoia, including visualization engines, archival storage movers, BlueGene/L, Purple, and TLCC07 Linux clusters.

Planned activities in FY09:

- Award Sequoia platform build vendor contract
- Award Sequoia platform R&D vendor contract
- Hold review(s) for Sequoia planning
- Plan for Sequoia ID (Dawn) platform delivery

Expected deliverables in FY09:

- Sequoia Performance Baseline/Construction Readiness CD-2/3 package
- Sequoia ID (Dawn) demonstration
- Sequoia ID (Dawn) early science runs
- Sequoia ID (Dawn) system transition to the Secure Computing Facility (SCF)

Preliminary planned activities in FY10:

- Prepare for Sequoia build Go/NoGo
- Plan for Sequoia parts commit and options
- Plan for Sequoia demo and early science runs

WBS 1.5.4.1-LANL-001 Systems Requirements and Planning

The Systems Requirements and Planning 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 is to define requirements and potential system architectures for future capability platforms that meet ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts and provides project management support for those efforts. In FY09, this project will focus on the project management of the hybrid Roadrunner Phase 3 system to be deployed at LANL. This includes overall project planning for the delivery, acceptance, deployment, and system integration of the Roadrunner Phase 3 final system. This section maps to the Roadrunner Project Element "3.1 Project Management."

The primary capability is to ensure that the FY09 Level 2 Roadrunner milestone for system integration readiness is accomplished by meeting the completion criteria in the milestone description.

Planned activities in FY09:

- Provide the overall project planning for the delivery and deployment of the Roadrunner phase 3 system, including system integration and management
- Continually use the project execution model process for the procurement and integration of the Roadrunner system
- Address any issues with the petascale computing environment through continued tri-lab planning

Expected deliverables in FY09:

- Roadrunner Phase 3 final system delivered, acceptance testing completed, and initial deployment started including network integration

Preliminary planned activities in FY10:

- Complete CD4 for Roadrunner Phase 3 to transition to operations

WBS 1.5.4.1-SNL-001 Red Storm Capability Computing Platform

Red Storm is a tightly coupled massively parallel processor compute platform with approximately 125 teraFLOPS of peak processing capability. The machine uses 2.4 GHz dual-core AMD Opteron processors and a custom, very high performance, 3D-mesh

communication network. Red Storm has a total of 13,600 dual-core Opteron processors and over 38 terabytes of memory and approximately 1.9 petabytes of high-performance local disk split between classified and unclassified use. Red Storm produced a 102.2 teraFLOPS result on the latest HPL benchmark, which is slightly more than 80 percent of its theoretical peak performance. Cray now has over 20 sites and has 38 systems sold based on the Red Storm architecture. Red Storm is to be upgraded to approximately 284 teraFLOPS and with 75 terabytes of compute node memory by the end of FY08. Specifications and an expanded description are available at http://www.sandia.gov/ASC/redstorm.

Planned activities in FY09:

- Provide maintenance for Red Storm to ensure it remains a useful and stable platform for ASC applications.
- Run benchmark problems and test suites to quantify the practical peak performance of the upgraded Red Storm by determining the actual throughput gain of an NNSA application.

Expected deliverables in FY09:

- Full production availability of the upgraded Red Storm
- The quantified actual throughput gain of an NNSA application using the upgraded Red Storm quad-core processors

Preliminary planned activities in FY10:

- Continue production computing

WBS 1.5.4.1-LANL/SNL-001 Alliance for Computing at Extreme Scale Zia Capability Computing Platform

ACES is a joint collaboration between LANL and SNL defined under a Memorandum of Understanding for developing requirements and system architecture for ASC capability systems requirements definition, architecture design, procurement, key technology development, systems deployment, operations and user support. A joint design team will develop requirements and architectural specifications for the Zia capability system. The architecture and design of Zia will be optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs.

Planned activities in FY09:

- Complete the selection of the system architecture for the Zia system
- Utilize the NNSA Chief Information Officer Project Execution Model for the procurement and integration of the Zia platform
- Evaluate vendor proposals based on system requirements to meet the mission need identified in the Critical Decision

Expected deliverables in FY09:

- Contract for the Zia system awarded
- Partial delivery of the Zia system by September 2009

Preliminary planned activities in FY10:

- Complete the delivery of the Zia system
- Complete the system integration of the Zia system

- Complete Zia system acceptance testing, including the demonstration of key NNSA applications running on the full scale of the Zia platform

WBS 1.5.4.1-LANL/SNL-002 Alliance for Computing at Extreme Scale Architecture Office

The primary objective for the ACES architecture office is to define requirements and potential system architectures for future capability platforms that meet ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts and provides project management support for those efforts.

A thoughtful and systematic approach is required to create the design for systems at extreme scale. The ACES architecture office will coalesce mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends in the design process. When vendor designs/proposals are available, for example, in response to the ZIA request for proposal, they require equivalent analysis to identify and select a proposal that matches pre-established design criteria. The ACES architecture office will also identify in collaboration with the computer industry, critical technology gaps for future production capability systems. Technology development projects will be initiated to address these issues.

Planned activities in FY09:

- Integrate technology trends into the design of upcoming ASC capability systems
- Evaluate vendor offerings for compatibility with ASC Capability computing requirements
- Initiate technology development projects in support of future ASC capability platforms

Expected deliverables in FY09 include:

- Completed specification for next capability system to address anticipated workload requirements
- Identification of suitable offering for next capability system

Preliminary planned activities in FY10:

- Begin analysis for follow-on capability system

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 FY09

- Complete deployment of all TLCC07 systems:
  - At LLNL: Juno (a cluster of 8 scalable units (SU)), Eos (2 SUs), Hera (6 SUs, for of which are ASC's)
  - At LANL: Lobo (2 SU) and Hurricane (2 SU)

  - At SNL: Glory (2 SU), Whitney (2 SU), and Unity (2 SU)
- Integrate Tripod selected stack and tools deployed on TLCC clusters

WBS 1.5.4.2-LLNL-001 Tri-Lab Linux Capacity Cluster

The TLCC07 procurement was designed to maximize the purchasing power of the ASC program by combining all tri-lab ASC capacity computing procurements into a single procurement. It is anticipated that purchasing a huge set of common hardware components will lead to lower cost through high-volume purchases. In addition, by deploying a common hardware environment multiple times at all three sites, it is anticipated that time and cost associated with any one cluster will be greatly reduced. The common TLCC07 hardware environment will also accelerate the adoption of a common tri-lab software environment. This project's budget covers cost of procuring the TLCC07 systems. Efforts to support activities and deliverables come from other product areas within CSSE (for example, System Software & Tools and Common Computing Environment).

The TLCC07 procurement provided the ASC community with a large amount of capacity computing resources in late FY08. Integration of these clusters was delayed four months because of a manufacturing problem in the Barcelona processor. In FY09, we will finish planned deployment of all the tri-lab clusters including full integration on classified networks. We will then fully realize our strategy for quickly building, fielding and integrating many Linux clusters of various sizes into classified and unclassified production service through the concept of SUs. The programmatic objective is to dramatically reduce overall total cost of ownership of these "capacity" systems relative to best practices in Linux cluster deployments today. This objective strives to quickly make these systems robust, useful production clusters under the coming crushing load of ASC scientific simulation capacity workloads.

Planned activities in FY09:

- Approve security plan and update for Juno system on classified network
- Integrate Juno 8-SU cluster into the classified LLNL simulation environment and complete the benchmark and application acceptance tests
- Approve security plan and update for Eos system on classified network
- Integrate Eos 2-SU cluster into the classified LLNL simulation environment and complete the benchmark and application acceptance tests

Expected deliverables in FY09:

- Full acceptance and production deployment of Juno 8 SUs running TOSS on the classified network
- Full acceptance and production deployment of Eos 2 SUs running TOSS on the classified network

Preliminary planned activities in FY10:

- Explore a follow-on TLCC capacity procurement in FY11–FY12, if directed to do so by NNSA ASC Headquarters (HQ); this procurement would be handled by the TLCC technical team

WBS 1.5.4.2-LANL-001 Capacity System Integration

The Capacity System Integration project will continue system integration for all capacity systems at LANL. In FY09, the main focus is integrating the TLCC systems delivered to LANL in Q3 FY08, and continuing to support other capacity systems at LANL, including ongoing integration and testing of the Roadrunner Base system. The TLCC capacity systems will deliver 80 teraFLOPS total, in accordance with the request for proposal describing common system components located at each of the three weapons labs. The system integration effort will address complex integration issues relating to high-end terascale computing environments and provide the necessary resources to stand up such systems for production computing work.

Planned activities in FY09:

- Coordinate, integrate, and test the Roadrunner Base System in full production and GA
- Plan and deploy TLCC systems at LANL: Lobo is targeted for the open Turquoise network, and Hurricane will be in the secure partition

Expected deliverables in FY09:

- Increased production maturity of the Roadrunner Base System
- Full acceptance and production deployment at LANL of initial TLCC SUs
- Tripod selected stack and tools deployed on TLCC clusters (with Application Readiness Team) are integrated
- Applications Readiness Team stabilization and continued productivity of both TLCC clusters

Preliminary planned activities in FY10:

- Explore follow-on TLCC capacity procurements upon the direction of NNSA ASC HQ

WBS 1.5.4.2-SNL-001 ASC Capacity Systems

The purpose of the ASC Capacity Systems project is to support the acquisition, delivery and installation of new ASC capacity systems. Each of the three Sandia TLCC platforms are 2 SU in size. The initial system Unity was delivered to Sandia, New Mexico in June 2008. Whitney was delivered to Sandia, California in August 2008. Glory was delivered in September 2008 to New Mexico. All systems will be tested on the restricted network before being placed into production on the classified network in FY09.

The project is supported by analysis of SNL's portfolio of application needs for capacity computing systems within the context of the higher integrated ASC platform strategy of capability, capacity, and advanced systems. Efforts include definition of requirements for TLCC system procurements and collaboration with WBS 1.5.4.7 Common Computing Environment with respect to a common software stack for new and existing capacity systems.

Planned activities in FY09:

- Complete site integration and acceptance testing of Unity, Whitney, and Glory clusters.
- Complete Security Plans for Unity, Whitney, and Glory for placement into the Sandia Classified Network.

Expected deliverables in FY09:

- TLCC07 systems operating in SNL's production computing environment running the TOSS stack

Preliminary planned activities in FY10 include:

- Continue maintenance and operation of TLCC07 systems in SNL's production computing environment running the TOSS stack
- Support TLCC11 procurement process for the next generation of ASC capacity computing systems

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 test bed 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 FY09

- BlueGene/Q central processing unit ASIC and network ASIC design manufacturing release
- The Roadrunner system completion in early FY09 to achieve petaFLOPS-level computing performance to meet ASC weapons computing requirements
- Acceptance testing of Roadrunner Phase 3 system
- Second IAA workshop on Memory Opportunities for High-Performance Computing
- Parallel version of SST
- Quantify the ability to predict failures or provide requirements to enable prediction on extreme-scale systems
- Optimized MPI-IO, parallel HDF5, and NetCDF libraries on ASC platforms; testing support for end-to-end I/O optimization
- Demonstrate 1GB/s, n+3 parity generation/check using GPGPU acceleration for RAID6

WBS 1.5.4.3-LLNL-001 BlueGene/P and BlueGene/Q Research and Development

The BlueGene/P and BlueGene/Q Research and Development project is a multi-year NNSA and Office of Science R&D partnership with IBM on advanced systems. It targets the development and demonstration of hardware and software technologies for 1-petaFLOPS and 10-petaFLOPS systems. The BlueGene/P hardware is based on an extension of the highly successful BlueGene/L architecture with more cores per node, 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, the E&TS division, and Argonne National Laboratory and the ASC tri-labs. In FY08, a BlueGene/P system was delivered to Argonne. Follow-on BlueGene/Q system design targets a 20-petaFLOPS system at the end of the contract.

This project incorporates requirements from the DOE laboratories, especially Argonne and LLNL, to have input into design choices and system testing for microprocessors, node architectures, and interconnects. The DOE laboratories also provide critical input on software, ensuring appropriate capability and features for the design target.

Planned activities in FY09:

- Review the BlueGene/Q compiler, packaging and prototype
- Simulate application kernels to expose thread-parallel performance aspects
- Simulate application use of key features of BlueGene/Q including transactional memory
- Investigate other code improvement opportunities

Expected deliverables in FY09:

- Manufacturing release of central processing unit ASIC design, network ASIC design, and system packaging

Preliminary planned activities in FY10:

- Review the system design for BlueGene/Q

WBS 1.5.4.3-LLNL-002 Petascale Application Enablement

The Petascale Application Enablement project enables advanced application work to develop benchmarks for new platforms, such as Sequoia, and to adapt current codes to the expected new architectures. A primary target of this project is investigating ways to improve application thread performance for future many-core platforms. The project team efforts include both direct application work and benchmark development and testing.

Planned activities in FY09:

- Continue vendor interactions with respect to Sequoia application performance requirements
- Start initial science and weapons code testing on Sequoia ID (Dawn) system, especially measuring single node thread performance
- Investigate opportunities for thread-parallel performance in production applications

Expected deliverables in FY09:

- Evaluation of Sequoia benchmark results submitted in response to RFP

Preliminary planned activities in FY10:

- Continue science and weapons code testing on Sequoia ID (Dawn) system
- Continue with focus on code improvement opportunities

WBS 1.5.4.3-LANL-002 Roadrunner Phase 3 Delivery and Acceptance

The Roadrunner Phase 3 Delivery and Acceptance project maps to the Roadrunner WBS Project Element "3.2 Acquisition." The budget dollars in this element are platform dollars only that are associated with payments to IBM for the system contract or the analyst support.

The Roadrunner Phase 3 System delivery and acceptance will be completed. The final system is configured with hybrid nodes based on a hybrid architecture using IBM System AMD Opteron-based processors accelerated with IBM's Cell Broadband Engine ("Cell BE") blades.

Hybrid computing architectures are an important direction for HPC. The projected peak performance of the final system is over 1 petaFLOPS.

Planned activities in FY09:

- Complete delivery of system
- Acceptance test Roadrunner Phase 3 system

Expected deliverables in FY09:

- Final Roadrunner Phase 3 Final System installed
- Completing the acceptance test and taking ownership of the Roadrunner system

Preliminary planned activities in FY10:

None. This project is completed in FY09.

WBS 1.5.4.3-SNL-001 Advanced Systems Technology Research and Development

The Advanced Systems Technology Research and Development project will collaborate with the recently established Institute for Advanced Architectures and Algorithms (IAA) to help overcome some of the bottlenecks that limit supercomputer scalability and performance through architectures and software research. For the current FY, the architectures focus will be on 1) advanced memories, 2) high speed interconnects and 3) power management techniques to reduce runtime power consumption of current and future platforms.

The project will research and evaluate the quantitative impact of five classes of advanced memory operations on NNSA/ASC applications: 1) memory controller operations, 2) basic buffer-oriented "in-memory" operations, 3) atomic memory operations, 4) scatter/gather operations, and 5) enhanced caching and translation "look aside" buffer operations.

In addition, the project will develop new network interface controller architectures to enable efficient data movement from host memory to the high-speed interconnect and use simulation tools to better our understanding of the high-speed interconnect requirements of the NNSA/ASC applications.

Sandia has developed and deployed on Red Storm a method to place unused cores in a power-saving state, which has been estimated to be saving several $10Ks in power costs for the platform over a year's time. This FY, the project will better quantify the use of power by important NNSA/ASC applications and research methods to improve performance/power ratios in future platforms. Additionally, this project will focus on

developing very-low power new processor designs that will allow dramatically more efficient processing for algorithms and applications on the path to Exascale computing.

The project will also explore software capabilities that will be essential for increasing the performance and scalability of applications beyond that provided by general-purpose operating systems, while providing the functionality necessary to deal with evolving compute node architectures, networks, parallel programming models, and applications. The overall goals of this project are to increase application performance on future many-core processors, ease the transition of applications to alternative programming models, and provide robust system software support for scaling applications to a million MPI tasks. This will include work on enabling the user to interact easily with the data in order to perform analysis.

Planned activities in FY09:

- Work with industry and laboratory partners to develop alternative memory architectures
- Develop a new high-speed interconnection controller to increase data movement efficiencies
- Complete the chip multiprocessors milestone
- Investigate techniques for managing power consumption and increase performance/power consumption ratios in future platforms

Expected deliverables in FY09 include:

- Second workshop on Memory Opportunities for High-Performance Computing. Final report will document the outcomes of the workshop.

Preliminary planned activities in FY10:

- Complete the advanced memory subsystems milestone
- Instantiate new research directions, in cooperation with the IAA, to impact the future capability platform

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.

System Software and Tools Deliverables for FY09

- Deployment of TOSS 1.2 (based on RHEL 5.3)
- Deployment of OpenFabric Enterprise Distribution 1.3
- Moab and SLURM ported to Sequoia ID (Dawn)
- Certification of Dawn application development tools

- Certification of TLCC Tripod application development tools
- Reports on performance modeling and measurement verification supporting the Roadrunner Base system production operations, Roadrunner Phase 3 build up and integration, and future systems such as Zia and Sequoia
- A suite of tools to assist in diagnosing performance problems, including a refinement of P-SNAP (the Performance and Architecture Laboratory System Noise Activity Program), anomaly detection, and deficiencies in network components
- Defined performance targets for FY09 releases for a code project
- Improved metrics for hybrid computer systems
- Architectural recommendations for code refactorization efforts
- Performance analysis tools integrated into a system health monitoring test suite
- Runs completed to accept Roadrunner Phase 3 system
- System component models for resilience on the cell architecture
- Tools for performance improvements on the cell architecture
- An infrastructure to efficiently use the memory of the Roadrunner system within the Crestone project
- Analysis of the performance and scalability impact of introducing new functionality into light weight kernels and high speed communication protocols
- Analyses of the performance of several ASC applications on multi-core processors, with identification of key bottlenecks

WBS 1.5.4.4-LLNL-001 System Software Environment for Scalable Systems

The System Software Environment for Scalable Systems project provides system software components for all the major platforms at LLNL, 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 system software components needed to augment Linux and required proprietary operating systems that function in a manageable, secure, and scalable fashion needed for LLNL ASC platforms.

This project includes work on developing, modifying, and packaging the TOSS, and developing scalable system management tools to support the operating system and interconnect (for example, TOSS and InfiniBand monitoring tools), as well as the resource management environment (Moab and SLURM) to queue and schedule code runs across LLNL systems.

Planned activities in FY09:

- Continue ongoing development and support TOSS software
- Continue planning for CHAOS 5/TOSS 2.0 software releases
- Continue ongoing development and support of Moab and SLURM

Expected deliverables in FY09:

- Deployment of TOSS 1.2 (Based on RHEL 5.3)

- Deployment of OpenFabric Enterprise Distribution 1.3
- Moab and SLURM ported to Sequoia ID (Dawn) system

Preliminary planned activities in FY10:

- Continue ongoing TOSS software development and support
- Deploy TOSS 1.3 (based on RHEL 5.4)
- Develop beta release of CHAOS 5/TOSS 2.0 contingent on RHEL 6 schedule
- Develop InfiniBand SAN support for Sequoia
- Continue ongoing development and support of Moab and SLURM
- Port Moab and SLURM to Sequoia system

WBS 1.5.4.4-LLNL-002 Applications Development Environment and Performance Team

The Applications Development Environment and Performance Team project provides the code development environment for all major LLNL platforms, supports user and code productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools such as platform partners, independent software vendors, and the open source community. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The project resolves bug and user trouble reports, including interactions with the software providers to fix problems.

The elements of the development environment 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.

Interactions between project members and code development teams ensure high performance use of existing systems and supports customer-based planning of future improvements to the environment. Similarly, long-term relationships with external partners, such as IBM, TotalView Technologies, the Krell Institute and Openworks, ensure that project members can resolve trouble reports quickly and avoid unnecessary duplication of existing capabilities.

Planned activities in FY09:

- Provide and maintain Purple and BlueGene/L code development environments
- Coordinate integrated design code scaling for Sequoia
- Provide and refine the common tri-lab environment for TLCC
- Provide and refine the development environment for Sequoia ID (Dawn) system
- Track and refine petascale development environment approaches
- Develop new techniques to improve robustness and performance of ASC codes
- Interact with the ASC code teams and vendors to improve software products

Expected deliverables in FY09:

- Certification of the Sequoia ID (Dawn) system applications development tools environment

- Certification of the TLCC Tripod applications development tools environment
- Continued deployment of highly scalable code correctness tool suite

Preliminary planned activities in FY10:

- Continue code development environment support on all LLNL ASC platforms
- Deploy full production version of highly scalable code correctness tool suite
- Identify and develop refinements of the code development environment for existing and future capacity and petascale systems
- Explore nested node concurrency programming model interfaces and performance
- Continue to support users and interact with vendors to serve user needs

WBS 1.5.4.4-LANL-001 Roadrunner Computer Science

The Roadrunner Computer Science project supports the success of the full Roadrunner system, from enhancing production maturity of the base system through the support of the hybrid system running at a sustained petaFLOPS. This project performs work in the areas of systems, communications, performance measurement, analysis and modeling, tools, and architecture for the Roadrunner project during the initial hardware availability, build-up, and pre- and post-installation of the system. This project maps to the Roadrunner Project Element "3.5 Software Tools/Programming Models."

The project will provide a set of essential capabilities leading to a successful Roadrunner system for applications of interest. The tools to be developed mainly target memory and communication, in addition to the acceptance testing from a performance perspective for the system.

Planned activities in FY09:

- Execute an acceptance test of Roadrunner performance
- Complete performance analysis and modeling of Roadrunner (Level 2 ASC milestone)

Expected deliverables in FY09:

- Reports on performance modeling and measurement verification supporting the Roadrunner base system production operations and full Roadrunner Phase 3 hybrid system build up and integration
- Report on performance acceptance testing

Preliminary planned activities in FY10:

None. This project is completed in FY09.

WBS 1.5.4.4-LANL-002 Software Technologies for Next Generation Platforms

The scope of the Software Technologies for Next Generation Platforms project is to measurably improve the usability, performance, reliability, efficiency, and productivity of petaFLOPS supercomputers, beyond the Roadrunner project, for nuclear weapons applications. The work will help focus the architectural choices and the software environment for next-generation capability machines, while having direct applicability to machines currently in use (Roadrunner). This work will focus on a radiation-

hydrodynamics workload of interest to ASC. This project will provide direct guidance to code developers and system vendors, including the development of hardware and software prototypes for the future. The work will be tightly integrated with the Roadrunner computer science project, the Roadrunner weapons science project, the Roadrunner algorithms project, the code performance and throughput project, the accuracy project, and the integrated codes project.

Planned activities in FY09:

- Evaluate and examine achievable system performance on available state-of-the-art multi-core processors, multi-socket nodes, and systems (as available) through modeling.
- Examine the achievable performance and portability on multi-core processors using specific optimizations as well as the impact of using abstractions on performance.
- Develop a suite of system performance health tools by building on the by the PAL team on optimizing various large-systems; this work will ensure that, at any point in time, a system is achieving the highest possible performance level.
- Continue implementation of fault-tolerant domain-specific language (DSL) compute-intensive kernels.
- Optimize system software architecture for increased performance and operating system noise reduction; expand memory tool beyond static allocation and analysis to dynamic allocation and runtime analysis.
- Determine requirements for multi-core/multi-processor systems.
- Compiler techniques for emerging heterogeneous architectures: Examine and prototype the impact heterogeneous, multi-core architectures will have on existing compiler infrastructure. In particular, data-flow analysis and scheduling algorithms will be explored. This work will leverage the open-source LLVM project (www.llvm.org) and provide the foundation and various building blocks for the implementation of domain specific languages.

Expected deliverables in FY09:

- Application of prefetching techniques developed in FY08 to hide memory latency costs to specific applications including transport and hydro
- A suite of tools to assist in diagnosing performance problems, including a refinement of P-SNAP (the PAL System Noise Activity Program), anomaly detection, and deficiencies in network components (the aim will be to provide a usable suite of tools that could be built into the normal diagnostics of a system and run nightly, or weekly, or on reboot, or even on job launch)
- Enhancement and expansion of fault-tolerant Direct Numerical Simulation (DNS) prototype developed in FY08
- Implementation of DSL designed in FY08 for memory extraction, layout, and communication
- Reports on performance of existing and future systems, including Zia and Sequoia
- Memory tool for dynamic allocation and run-time analysis
- LLVM framework expansion to provide support for dealing with hierarchical memory systems and heterogeneous cores using select accelerators (for example, IBM CBE and GPUs); this will include scheduling based on memory locality,

arithmetic complexity, and suitability of a particular code for a given processor architecture; a prototype compiler will be completed by year end

Preliminary planned activities in FY10:

- Provide the ASC community with a suite of memory and communication tools that support the program's application development efforts, assist in developing prototype application software, and enable more accurate system architecture designs

WBS 1.5.4.4-LANL-003 Code Performance and Throughput

The goals of the Code Performance and Throughput project are to measurably improve the performance of the integrated X Division weapons code projects at each major release, provide predictive performance tools for the same projects, and inform the code re-factorization project.

This project will develop performance prediction tools for the tri-lab user community while meeting specific requirements for the hybrid Roadrunner system. To ensure relevance, this project will deploy staff directly on code projects to enhance code performance in each major release. As an overarching goal the project will seek to estimate the runtime based on large-scale variables like physics method, computer system, data sets, and other factors. This project will define targets for enhanced performance as a function of time, deploy computer science expertise to both major X Division weapons codes projects, and enhance the performance of the codes against those targets as delivered to the user community.

Planned activities in FY09:

- Continue deploying performance metric-tracking software on tri-lab systems and investigate working with ASC code teams at other labs
- Develop support for hybrid computer systems (for example, Roadrunner) addressing issues of exotic ISA hardware counters and asynchronous data movement.
- Complete web-based accessible performance data for a range of X Division (weapons codes) and CCS Division (transport) code projects
- Assist ASC physics code teams to use the performance metric-tracking system and other tools to identify performance bottlenecks in current production codes with a goal of increasing code performance at each major release
- Deploy a prototype predictive system based on collected data and heuristics, to aid end-users and managers in planning for effective use of ASC multi-physics codes and resources
- Investigate cross-fertilization of analytic performance modeling with the prototype system
- Work with ASC physics code teams to inform, from a computer-science perspective, their on-going refactorization efforts for current and future ASC platforms

Expected deliverables in FY09:

- Defined performance targets for FY09 releases for a particular X Division code project
- Improved metrics for hybrid computer systems

- Plan of attack on major code hotspots for a particular X Division code project
- Prototype toolset for physics-based performance prediction for an X Division code project
- Architectural recommendations for code refactorization efforts

Preliminary planned activities in FY10:

- Continue ongoing improvements in metrics
- Define performance targets for all X Division code projects
- Complete performance predictors for an X Division code project
- Study software architectural for code refactorization efforts

WBS 1.5.4.4-LANL-005 Software Support

The Software Support project works to establish a strong parallel development and analysis tool capability for current and next generation HPC platforms. It is focused on working with the HPC tool community to identify, plan, and integrate tools into production environments and establish a solid support structure for sustainability. It includes supporting the development and implementation of tools based on a tools strategic plan. The plan includes cross-laboratory partnerships and external collaborations that focus on performance tools required for programming model support.

Capabilities include tool strategic plan development based on current and next generation platform planning, integration with HPC community tool development efforts, and tool development and production integration capability.

Planned activities in FY09:

- Continue building Open MPI support capability by engaging community support model
- Continue Open|Speedshop (O|SS) integration in application analysis and system scalability
- Support Zia platform software planning
- Continue involvement with DOE NNSA and Office of Science laboratories in petascale tool planning and development programs

Expected deliverables in FY09:

- Integrate performance analysis tools to a system health monitoring test suite
- Create a tri-lab consortium for maintenance-level support for open source software used primarily by ASC

Preliminary planned activities in FY10:

- Continue to drive OpenMPI community efforts and optimization efforts
- Continue efforts with O|SS toward a scalable infrastructure and additions of plugin tools to support multicore architectures

WBS 1.5.4.4-LANL-006 Applications Readiness

The Applications Readiness project addresses issues with getting applications running on current and incoming computing systems at LANL. Working with subsystem teams such as systems management, file systems and I/O, archive, and tools, the Applications Readiness team identifies causes of system failures or unexpected behavior and deploys fixes in production. The project goal is that system users are able to make productive use of the systems with their applications to solve their problems.

The project provides drill down and squash problems (create small problem reproducers, identify cause, work with the relevant vertical(s) to find solution, and verify deployed solution), periodic stress testing/regression of production machines, new software version regression testing, verify system configuration and software stack for new or upgraded systems with real user applications, and metrics and analysis/profiling.

Planned activities in FY09:

- Continue to take on the hardest and most elusive problems with existing machines including the Roadrunner Base capacity system
- Ensure an orderly application deployment on Roadrunner phase 3 and TLCC systems
- Assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner phase 3 system
- Assist system management personnel with problem investigation and recreation

Expected deliverables in FY09:

- Runs completed enhance production maturity of Roadrunner Base system
- Runs completed to accept TLCC systems into production
- Runs completed to transition Roadrunner Phase 3 system to an operational state
- Report on characterization of memory reference patterns for one or more applications of interest

Preliminary planned activities in FY10:

- Continue to take on the hardest and most elusive problems, or newly identified problems, with existing production machines
- Ensure an orderly application deployment on the Zia system
- Continue to assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner phase 3 system
- Assist system management personnel with problem investigation and resolution

WBS 1.5.4.4-LANL-007 Productivity Project

The Productivity Project provides direct support to LANL ASC code projects for efficient usage of ASC machines. This includes application performance, resilience, scheduling, I/O, and data structures. The project demonstrates optimization procedures based on model output and runs integrated simulations to demonstrate optimization productivity improvements. The project identifies root causes for problems related to I/O and memory usage. The project develops strategies and plans to improve the

performance of ASC code projects. The project focus is on ASC code projects to improve their productivity.

Planned activities in FY09:

- Based on a fundamental understanding of computer architecture, its memory management, and the physics in the code, systematically design new data structures for Crestone project, which include a more efficient AMR strategy
- Develop system component models for resilience on computers with cell architectures
- Gather and analyze data for system models
- Develop optimization procedure for cell architectures through data gathering, monitoring, and scheduling
- Demonstrate the performance/productivity improvements through real applications on cell architectures
- Identify the footprint of memory usage within Crestone project
- Significantly modify various modules within Crestone project so that large problems with many materials could be run due to the improvement of the memory usage

Expected deliverables in FY09:

- System component models for resilience on the cell architecture
- Tools for performance improvements on the cell architecture
- An infrastructure to efficiently use the Roadrunner memory within Crestone project

Preliminary planned activities in FY10:

- Investigate additional optimization procedures on petaFLOPS machines
- Develop tools to optimize production runs on petaFLOPS machines
- Develop tools to help generation and partition for large 3D unstructured meshes for ASC code projects
- Develop and implement a plan for the connection between ASC code projects

WBS 1.5.4.4-SNL-001 Software and Tools for Scalability and Reliability Performance

The Software and Tools for Scalability and Reliability Performance project supports software R&D to address scalability and reliability performance of future computational systems. This project has multiple components with the unifying theme of continuous improvements in the software infrastructure supporting ASC applications. The end goal is efficient, accurate, and timely results to application runs done in support of the ASC mission. We employ multiple, but focused, approaches to achieve this end goal. Highly scalable system software is critical and must evolve for future microprocessor architectures, such as multi-core. Evolutionary system software change is important for easing the transition to new hardware for the applications. Thus, highly scalable MPI will be complemented with new high-speed communication algorithms for non-message-passing techniques, such as on-chip threading and shared memory algorithms. Time to completion can be minimized if system reliability can be improved through software resiliency techniques. Development of these features must be complemented with performance analysis on ASC applications. Activities in this task are prioritized by

their applicability to the upcoming Zia platform. Additionally, we will investigate the coupling between OS services, hardware and applications behavior to develop systems that allow highly scalable behavior without manual intervention.

Planned activities in FY09:

- Develop techniques using lightweight kernel software to support, in a highly-scalable manner, emerging requirements for such features as dynamic libraries, lightweight threading, and multi/hetero processor architectures

Expected deliverables in FY09 include:

- Analysis of the performance and scalability impact of introducing new functionality into light weight kernels and high speed communication protocols
- Analyses of the performance of several ASC applications on multi-core processors and identify key bottlenecks
- Prototype new framework design for resource monitoring, task (re-)direction, resource redirection in multicore architectures

Preliminary planned activities in FY10:

- Provide more evolutionary system software support for changes in HPC hardware
- Execute runtime and scalability performance studies of the Zia platform

WBS 1.5.4.4-SNL-003 System Simulation and Computer Science

The System Simulation and Computer Science project will work with the recently established Institute for Advanced Architectures and Algorithms and will focus on: 1) software and methods for multi-core microprocessors and extreme scalability and 2) software and methods for ensuring adequate resilience and fault-tolerance characteristics of extreme-scale platforms.

Given the extreme cost of deploying a capability machine, as well as the high cost of developing the complex multi-physics codes that run on them, it is important to use a systematic design approach that permits a design informed by predictions of application performance on these machines. This project, in collaboration with the IAA algorithms and computer science effort, will develop a multi-scale simulation capability leveraging the prior investment in SNL developed tools: the Structural Simulation Toolkit (SST) (micro-level), Seshat (meso-level) and ArchSim (macro-level).

Resilience and fault-tolerance are key issues for current and future HPC systems. The mean time between interrupts approximately scales with the inverse of the number of parts. For Red Storm, the mean time between interrupts on the full machine is approximately 14 hours. Zia may have 4 times as many processors as Red Storm and a future Exascale system will be a factor 100 to 1,000 times larger. This project will research and develop methods and solutions to ensure that future extreme-scale systems will have adequate resilience and hence provide a productive environment for users. Other important aspects of this work include developing an ability to predict failures and to determine root causes.

Planned activities in FY09:

- Research methods for increasing performance and scalability of applications on future architectures
- Develop a multi-scale simulation environment to support architecture performance analysis

- Research and develop methods for increasing the resilience of extreme-scale systems

Expected deliverables in FY09 include:

- Demonstration of a simple processor time estimator for macro-scale simulator, ArchSim
- Full integration of ns-3 network simulator into ArchSim
- Demonstration of usage of simulation to measure the effect of network topology changes on a Sandia application
- Deliver a parallel version of SST
- Provide a quantification capability to predict failures or provide requirements to enable prediction on extreme-scale systems

Preliminary planned activities in FY10:

- Refine multi-scale simulation environment with more accurate network models and better processor time estimators
- Expand application of simulators to more ASC applications

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

This level 4 product provides I/O (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 test bed 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.

Input/Output, Storage Systems, and Networking Deliverables for FY09

- Requirements, design documents, and initial coding for HPSS R8.1
- Production deployment of HPSS R7.1
- Production deployment of Lustre 1.8
- File system and archival storage hardware deployment to support Sequoia ID system (Dawn)
- Switching infrastructure hardware integration and deployment to support Dawn
- Deliver HPSS v7.2
- Provide new version of parallel storage interface (PSI)
- I/O infrastructure deployed to the ACES Zia system

- Support for the transition of Roadrunner Phase 3 from Open Science to secure production
- Report/presentation on I/O pattern and file system usage for the Roadrunner Phase 3 Open Science stability period
- Report/presentation on full I/O path fail over for global parallel file system (GPFS)
- Optimized MPI-IO, parallel HDF5, and NetCDF libraries on ASC platforms; testing support for end-to-end I/O optimization

WBS 1.5.4.5-LLNL-001 Archive Storage

The Archival Storage project provides end-to-end, long-term, high performance, archival storage services to ASC customers. This includes a collaborative software development effort between the tri-labs, Oak Ridge National Laboratory, and IBM, as well as deployment and support of archival storage software and interfaces for tri-lab ASC customers on unclassified and classified networks. It also includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and storage media and the ongoing technology refresh and data stewardship.

Archival storage system software (currently, HPSS) provides scalable, parallel archival storage interfaces and services to customers running at the tri-labs. HPSS distributes data across a configurable amount of storage units and removes other limits to scaling including number of files, directories, and concurrent users.

A world-class array of storage hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms and not hinder computation. This includes disk arrays, tape subsystems, mover nodes, storage-area-networks, networks, robotics and petabytes of media. Together, this hardware and software supports high-speed parallel transfer rates, currently in excess of 5 GB/sec. at LLNL, into an unlimited data store, at a current capacity of over 10 petabytes in a single name space.

Planned activities in FY09:

- Develop HPSS R8.1 through requirements, design, and initial code phases. HPSS R8.1 is expected to feature an architecture utilizing distributed core servers and partitioned metadata to meet the extreme scalability requirements of petascale computing in the Sequoia time frame.
- Deploy HPSS R7.1 in production environments, including upgrades to all requisite software infrastructure suites, including Kerberos, Lightweight Directory Access Protocol, and DB2.
- Migrate data from aging tape technologies to decommission robotic tape silos, drives, and media by end of CY2010, when support ends.
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support and maintenance of all archival storage hardware and media, customer and interface support, and ongoing tech refresh, and data stewardship of LLNL archives.

Expected deliverables in FY09:

- Requirements and design documents and initial coding for HPSS R8.1
- Production deployment of HPSS R7.1
- Hardware evaluations of core server and metadata subsystems for HPSS R8.1

- Archival storage hardware and software procured and deployed to support expansion of ASC platforms

Preliminary planned activities in FY10:

- Develop HPSS R8.1 through code review, system test, integration test phases, and beta release
- Plan for operations and prepare for deployment of HPSS R8.1

WBS 1.5.4.5-LLNL-002 Parallel and Network File Systems

The Parallel and Network File Systems (NFS) 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, the deployment and support of ubiquitous NAS services for home, project, and scratch space, and the I/O support of various programming interfaces for parallel I/O.

This project deploys and supports 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. It actively works with Sun Microsystems to add Lustre file system scalability and reliability enhancements required by TLCC, Dawn, and Sequoia platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

Planned activities in FY09:

- Provide Lustre file system performance and scalability enhancements in support of Lustre 1.6.5 and 1.8 releases
- Provide all necessary file system enhancements related to TLCC and Dawn platforms
- Work with Sun Microsystems to integrate Sun's Zettabyte Files System (ZFS) into Lustre targeting the Sequoia timeframe to provide needed scalability and reliability enhancements
- Investigate the ZFS file system as a possible backend use on NAS servers
- Maintain GPFS, NAS, and Lustre parallel file system support, middleware, and higher-level I/O library support for users
- Analyze and evaluate the security features of Sun's ZFS functioning as the backend of a NFSv4-based NAS server

Expected deliverables in FY09:

- Deployment of Lustre release 1.8 into production in Open Computing Facility (OCF) and SCF
- Deployment of Dawn file system hardware infrastructure
- Replacement of NAS servers in both OCF and SCF

Preliminary planned activities in FY10:

- Complete ZFS-based Lustre file system development and testing in preparation for Sequoia petascale environment
- Deploy NFSv4 as NAS server between tri-lab sites

WBS 1.5.4.5-LLNL-003 Networking and Test Beds

The Networking and Test Beds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect subsystems in support of current and future systems and environments. This work relies heavily on an adequately provisioned test bed, 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 subsystems.

Planned activities in FY09:

- Perform research and testing for technologies and products for interconnects, LANs and WANs and National Security Agency (NSA) Type 1 encryptors, file system servers, clients and disks, with special focus on emerging DataCenter Ethernet switches, additional features in InfiniBand, other interconnect and 10GigE related technologies, and 10GigE NSA Type 1 encryptors in support of future ASC petaFLOPS systems
- Study developing load balancing and multipath routing for previously listed networks
- Apply testing results to optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems

Expected deliverables in FY09:

- Installation and integration of switching infrastructure for Sequoia ID system (Dawn)

Preliminary planned activities in FY10:

- Continue to leverage tri-lab activities in I/O related hardware and software, and seek to improve the reliability, performance and manageability of the I/O subsystems in production
- Research and test to determine which technologies and products should be considered for insertion into production to meet the growing I/O performance and capacity requirements

WBS 1.5.4.5-LANL-001 File Systems and Input/Output Project

The File Systems and I/O 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 enterprise-wide 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.

This 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 500 GB/sec., global file system I/O rates beyond 500 GB/sec., and latencies in the 1-microsecond range. All this performance must be provided in an integrated, usable, reliable, and secure way. Data transfer and storage bottlenecks are now a critical concern for current-generation, HPC 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.

Planned activities in FY09:

- Continue work on deployment of I/O and storage infrastructure for Roadrunner Phase 3
- Assist users with I/O issues on Roadrunner phase 3 for open science runs with a special focus on weapons applications
- Assess and support deployment of the ACES Zia machine
- Continue to develop and tune full I/O path fail over for GPFS
- Demonstrate a single name space NFS service
- Test and demonstrate enhanced scalable metadata operation

Expected deliverables in FY09:

- I/O infrastructure deployed to the ACES Zia system
- Support for the transition of Roadrunner Phase 3 from Open Science to Secure production
- Report/presentation on I/O pattern and file system usage for the Roadrunner Phase 3 Open Science stability period
- Report/presentation on full I/O path fail over for GPFS

Preliminary planned activities in FY10:

- Support deployment of enhanced scalable metadata operation

WBS 1.5.4.5-LANL-002 Archival Storage Design and Development

The Archival Storage Design and Development project includes services for HPSS and PSI software development by LANL for the purpose of supporting ASC customers from LANL, LLNL, and SNL. 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 LANL colleagues on user interface issues and ensures that PSI functions with each new release of HPSS.

HPSS is software that manages petabytes of data on disk and robotic tape libraries. HPSS provides highly flexible and scalable hierarchical storage management that keeps

recently used data on disk and less recently used data on tape. HPSS uses cluster, local area network (LAN) and/or SAN technology to aggregate the capacity and performance of many computers, disks, and tape drives into a single virtual file system of exceptional size and versatility. This approach enables HPSS to easily meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored. HPSS provides a variety of user and file-system interfaces ranging from the ubiquitous vfs, ftp, samba and nfs to higher performance pftp, client API, local file mover and third party SAN (SAN3P). HPSS also provides hierarchical storage management services for IBM GPFS.

Planned activities in FY09:

- Develop requirements for 2009 release of PSI
- Design, develop and test 2009 release of PSI
- Finalize requirements for release 7.2 and 8.0 of HPSS
- Design, develop, test and release HPSS 7.2
- Design, develop and test HPSS 8.0
- Provide short term functional updates to the existing HPSS code base
- Provide second level support for archival storage deployment team

Expected deliverables in FY09:

- PSI requirements document
- PSI 2009 release
- Requirement Documents for HPS 7.2 and HPSS 8.0
- Release 7.2 of HPSS
- The FY10 release of HPSS 8.0

Preliminary planned activities in FY10:

- Continued Level 2 support for the production archive
- Initiate next-generation archive planning

WBS 1.5.4.5-SNL-001 Scalable Input/Output and Storage Systems

The Scalable I/O and Storage Systems project provides a scalable I/O infrastructure for ASC platforms in support of national security simulation needs and in direct alignment with ACES. In that capacity, the Scalable I/O project works with vendors and researchers to provide reliable, high-performance, easily used scalable I/O libraries and file systems that make optimum use of backend storage systems for petascale and beyond. It also educates users on how to ensure optimum I/O performance from their application, and plays an influential role in a number of standards bodies such as POSIX, parallel NFS, and NFS-RDMA.

Planned activities in FY09:

- Participate in ACES to help evaluate file system and I/O candidate solutions for Zia. In partnership with LANL personnel, work with application developers to optimize I/O performance on this new machine.

- Participate in ACES to help satisfy requirements and capabilities for the TLCC Trinity machine. Identify gaps, propose solutions, and work with vendors to address identified gaps in existing and potential future products for Trinity.
- Work with vendors on R&D to improve and adapt high-level I/O libraries such as netCDF, HDF5, and MPI-IO for the ACES platforms.
- Provide high-performance, reliable parallel file system I/O for existing ASC platforms, including the SYSIO library that provides virtual file system capability for Red Storm compute nodes.
- Participate in POSIX standards committee, parallel NFS and NFS/RDMA, Parallel Data Services Institute ties with NNSA, and NNSA obligations on the FAST-OS grant for IOFSL.

Expected deliverables in FY09:

- Optimized MPI-IO, parallel HDF5, and NetCDF libraries on ASC platforms; testing support for end-to-end I/O optimization
- Version 7 of HPSS operational on SNL storage platforms

Preliminary planned activities in FY10:

- Ongoing support of parallel file system I/O for ASC platforms
- Ongoing vendor support and guidance for R&D to improve high-level library support for ACES platforms

WBS 1.5.4.5-SNL-003 Archival Storage

The archival storage project represents SNL's participation in the DOE HPSS Consortium development project (described in detail in previous sections). HPSS provides the archival storage solution for ASC systems and is in direct alignment with ACES. SNL's role in the HPSS project is to collaborate with tri-lab developers to design, implement, and test solutions that meet ASC requirements for all three labs.

Planned activities in FY09:

- Develop and test of security, parallel file transport and performance improvements for current and prospective versions of HPSS
- Finalize 8.1 requirements (focuses on dynamically scaling core servers for a petascale environment)
- Ongoing support of PFTP/TA/LDAP/Security

Expected deliverables in FY09:

- Version 7.1 of HPSS

Preliminary planned activities in FY10:

- Support development and deployment of HPSS version 8
- Define requirements for next upgrade

WBS 1.5.4.5-SNL-004 Scalable Input/Output Research

The Scalable I/O project will drive I/O system enhancements for the next generation of extreme-scale systems through R&D in parallel file systems and I/O libraries. In particular, this project will explore the use of hardware accelerators to improve

bandwidth and reduce latency of I/O operations; investigate ways to exploit emerging storage architectures such as solid-state storage and Netezza systems to improve I/O; and research ways to use network resources to provide file-system caching and advanced in-transit data processing for I/O functions and high-level I/O libraries.

It is expected that most of the FY09 research and development of parallel file systems and I/O libraries will yield candidate technologies for future Exascale systems; these technologies, coupled with advances in network transport services, such as remote direct memory access (RDMA) and service guarantees, will create new opportunities to push state-of-the-art in I/O.

Planned activities in FY09:

- Research methods for improving parallel file system and I/O library performance through cooperative caching in the compute-node fabric
- Evaluate the use of emerging technologies (for example, solid-state storage, multi-core, Neteeza) for high-performance I/O and in-transit data processing
- Lead I/O research to improve performance and scalability of high-level I/O libraries such as netCDF, HDF5, and MPI-IO

Expected deliverables in FY09:

- Identification of issues and benefits associated with providing file-system support for cooperative caching on available compute nodes
- Demonstration of benefits of I/O library and file-system proxies.
- Optimized MPI-IO, parallel HDF5, and NetCDF libraries on ASC platforms; testing support for end-to-end I/O optimization
- Demonstration of 1GB/s, n+3 parity generation/check using GPGPU acceleration for RAID6

Preliminary planned activities in FY10:

- Develop I/O libraries that exploit acceleration technology for in-transit data processing
- Adapt parallel file system functionality to support caching in available compute nodes

WBS 1.5.4.6: Post-Processing Environments

This level 4 product provides integrated post-processing environments to support end-user 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 metadata and scientific data management, and general-purpose and application-specific visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data—the use of open-source foundations will continue to be an important strategy for development of shareable advanced techniques. The product must develop solutions to address interactivity, scaling and tri-lab access for petascale platforms, and data analysis techniques needed to support effective V&V and comparative analysis. Solutions for emerging platform architectures may in turn require customization and/or re-architecting of software to leverage hardware features. A continuing emphasis will be placed on tools for improving end-user productivity. The

product also provides and supports infrastructure including office and collaborative space visualization displays, mechanisms for image data delivery, and graphics rendering hardware.

Post-Processing Environments Deliverables for FY09

- Upgraded versions of movie player software and visualization resource management software, and enhanced remote rendering options
- Upgraded software for movie playing, resource management, remote rendering
- Upgraded PowerWall capabilities including stereo in additional theaters
- Version 2.3 of Hopper and Chopper file management tools
- Enhanced integrated distance and comparative/feature extraction visualization capabilities in software tools
- Implementation plan implementing accelerated visualization using cell or other technologies in user tool base
- High availability in the large visualization facilities
- Scientific visualization support for briefings, user visualizations, visualizations of a NIF inertial confinement fusion simulation, new weapons science visualizations for the Thermonuclear Burn Initiative program, and other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs
- PoP and GMV with modest updates to fix bugs and improve portability and maintainability
- New release(s) of ParaView that include enhancements for: petascale-sized simulation results; comparison of large scale simulated and experimental data; sensitivity analysis and UQ
- Level 2 milestone to provide scalable analysis tools for sensitivity analysis and UQ
- Preliminary implementation of latency tolerant visualization capability, in support of FY10 GA of Zia platform
- Delivery of relevant region detection tools, in support of ASC's UQ and V&V goals

WBS 1.5.4.6-LLNL-001 Scientific Visualization

The Scientific Visualization project provides high-end visualization and analysis capabilities, broken down into two primary activities: 1) research into new analysis, compression, and visualization techniques, and 2) operational support of the post-processing resources, including the visualization servers, displays, software, and facilities. Server efforts include planning, test bed prototypes, component and system testing, computer security and system deployment, system administration, and maintenance of visualization servers. Display efforts include high resolution/performance display devices for theaters and collaborative use areas. Operational support includes managing the use of 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 exploits the latest capabilities of clustering hardware, graphics processing unit advances, and parallel storage systems. Hardware capabilities include two

production visualization servers and several PowerWall clusters. The Linux visualization clusters are connected to a Lustre file system as well as a local NFS file system devoted to data analysis on the unclassified side. A video display infrastructure drives power walls and smaller displays. The project installs, maintain, and consults on a variety of software visualization tools, and supports demonstrations on the PowerWall. The project maintains both an unclassified and classified video production lab, each of which includes desktop systems with video editing, 3D modeling and animation tools, and various video peripherals, able to create DVDs and videotapes as needed or to support live demos.

Planned activities in FY09:

- Perform research and development in topics such as data compression, level of detail, and topology
- Participate in planning activities for petascale data analysis capabilities and provide ongoing hardware and software evaluation and technology watch
- Maintain and enhance the existing suite of tools and libraries developed by the visualization project in support of PowerWall usage
- Provide operational support for all visualization facilities, including supporting projection equipment, performing color alignment and projector alignment, and facilitating the use of the visualization clusters and associated storage
- Support ASC scientists through visualization and video support, including the support of PowerWall presentations, creating visuals and movies to high-list scientific data, and providing general user support

Expected deliverables in FY09:

- Upgraded versions of movie player software and visualization resource management software, and enhanced remote rendering options
- Upgraded PowerWall capabilities including stereo in additional theaters

Preliminary planned activities in FY10:

- Continue to maintain and enhance the PowerWall-related suite of software
- Perform research and development in topology, level of detail, and multi-core technology

WBS 1.5.4.6-LLNL-002 Scientific Data Management

As the amount of simulation data being generated grows, so does the difficulty of managing that data. The Scientific Data Management (SDM) project provides users with powerful and convenient ways to access, search, compare, and archive large-scale scientific data. This is achieved through the development of production-quality applications that enhance existing data management tools as well as provide new and innovative capabilities.

The Scientific Data Management tools team has expertise in data transport protocols, graphical user interfaces, web technologies, data representation, databases, and advanced system architectures. The team has decades of combined experience in designing and developing productivity-enhancing applications.

Hopper and Chopper are the principal products of this effort. Hopper is a Java-based file management tool that allows users to transfer and manipulate files and directories by means of a graphical user interface. Users can connect to and manage local and

remote resources using all file transfer protocols supported by ASC computing centers. Chopper is the command line version of the tool, useful in particular for automated, "background" file manipulations initiated from within applications.

Planned activities in FY09:

- Maintain and enhance the existing suite of tools and libraries developed by the scientific data management project
- Investigate inclusion of Hopper/Chopper in TOSS, and support use of our tools at other sites
- Explore additional petascale data management functions with users; for example, a crawler that scans scratch file systems for critical files, such as input decks, and automatically saves them to storage

Expected deliverables in FY09:

- Version 2.2 of Hopper and Chopper, featuring built-in file comparison capabilities and support for Livermore Computing's centralized storage access architecture
- Version 2.3 of Hopper and Chopper, featuring an enhanced directory viewer, support for direct-to-tape heuristics, improved scalability, and an expanded command launch mechanism

Preliminary planned activities in FY10:

- Continue to maintain and enhance the SDM suite of tools and libraries
- Explore the possible integration of the Hopper graphical user interface with a native desktop interface
- Gather user feedback and requirements for data management challenges faced in the petascale computing environment

WBS 1.5.4.6-LANL-001 Visualization and Insight for Petascale Simulations Project

The Visualization and Insight for Petascale Simulations Project is composed of two parts. The first part includes improved/automated data analysis capabilities, which encompasses ongoing work on comparative and quantitative visual analysis of ensembles of data. Automatically identifying features, measuring and comparing these features and using these features to prioritize work are key new capabilities that will be developed. The second part includes scaling to petascale interactive visualization and analysis, which encompasses identifying appropriate hardware resources to support petascale visualization, working on data reduction based software techniques such as intelligent data streaming to reduce the volume of data need to be moved off the petascale platform, and exploring the use of visualization kernels that run on the petascale platform (including data analysis, visualization, and rendering methods). Running on the petascale platform is important for two reasons: 1) to get around the latency limitations imposed by the architecture (for example, if you can analyze as the data is computed, you can avoiding writing to disk and then re-reading the data); and 2) for performance, using the power of the petascale platform to analysis petascale data.

The project team develops new visualization algorithms and systems to meet capability requirements for ASC petascale simulations. This work is required to address ASC workloads—massive data sizes, ensembles of results, and using unique supercomputing architectures.

Planned activities in FY09:

- Scaling to petascale interactive visualization and analysis, potentially accelerated using Roadrunner and Roadrunner-like technology
- Developing integrated distance and comparative/feature extraction visualization tool; automatic identified features will prioritize the data visualized by the distance visualization tool

Expected deliverables in FY09:

- Enhanced integrated distance and comparative/feature extraction visualization capabilities in software tools
- Implementation plan implementing accelerated visualization using cell or other technologies in user tool base
- Development of readers for Ensight and ParaView for the VPIC code

Planned activities in FY10:

- Advance petascale data analysis infrastructure by developing visualization, analysis and rendering software aligned with the new platforms, including the development of new approaches such as in-situ analysis and advanced feature extraction

WBS 1.5.4.6-LANL-002 Production Systems for Visualization and Insight Project

The primary goal of the Production Systems for Visualization and Insight Project is to provide world-class visualization infrastructure and support services in an efficient and cost-effective manner. This provides visualization services from the machine to the desktop for users in the ASC program. Technical staff 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.

This project also supports and maintains LANL's large visualization facilities, including the 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.

Capabilities include the support of large visualization systems, the support, maintenance and use of large visualization facilities, such as the CAVE and the PowerWall Theater, the production of custom visualizations for customers, and some code development.

Planned activities in FY09:

- Continue the support and maintenance of the large visualization facilities
- Support users in providing visualizations of simulations, experimental data, and engineering computer aided design (CAD) data on desktops and in facilities
- Specify contract management and requirements, including facilities, visualization cluster and EnSight contracts
- Plan visualization cluster upgrades

- Participate in Zia design and procurement activities, in particular, the integration of visualization capability into the Zia platform
- Analyze and further investigate data movement analysis project initiated in FY08, for the purpose of improved design of visualization platforms and perhaps visualization integration into capability and capacity platforms

Expected deliverables in FY09:

- High availability in the large visualization facilities
- User visualizations for reviews and briefings

Preliminary planned activities in FY10:

- Continue support and maintenance of the large visualization facilities
- Continue support of users in providing visualizations of simulations, experimental data, and engineering CAD data on desktops and in facilities
- Possible upgrade of Viewmaster rendering cluster

WBS 1.5.4.6-LANL-003 Physics-Based Simulation Analysis Project

The purpose of the Physics-Based Simulation Analysis Project is to help LANL weapons designers utilize the full power of the hardware and software infrastructure for visualization and data analysis developed and deployed by ASC to improve the physics understanding of their weapons simulations. To achieve this goal, the project has deployed within the design community in X Division a small group of individuals with expert knowledge in both visualization and weapons science to work directly with the designers. The job of this small group of experts is to help designers apply the full potential of the ASC visualization and analysis infrastructure to solve their analysis problems and to promote new weapons science discoveries using the ASC codes. The project also has a small effort to maintain and support PoP, an interactive 2D graphics post-processor, and General Mesh Viewer (GMV) and port these custom visualization and analysis codes to the new ASC platforms, such as the Roadrunner Base system.

In addition to working directly with the design community on its visualization and analysis problems, this group, in conjunction with the Production Visualization project, is responsible for all LANL activities related to the EnSight visualization and data analysis software. This includes maintaining the EnSight software installation lab-wide, providing local user support in the use of the software, and acting as a bridge between the LANL design community and the EnSight developers at Computational Engineering International (CEI) for problem reporting and resolution and for new feature requests. This group, with Production Visualization personnel, also directs all subcontracts LANL has with CEI related to new EnSight development and to onsite training and consulting.

Planned activities in FY09:

- Apply expert knowledge of weapons physics and the visualization software environment to iteratively develop solutions of interest to the program
- Work directly with designers in physics-based, iterative discovery process using petascale visualization and data analysis enabled tool (EnSight)
- Help support and maintain the EnSight software and direct its development activities under the new LANL EnSight development contract with CEI
- Document our work with joint publications co-authored with X Division designers on a variety of weapons science topics

- Maintain and support PoP and GMV and their associated link files, including GMV trials on Roadrunner with a Cell-based rendering library

Expected deliverables in FY09:

- Visualizations of a NIF inertial confinement fusion simulation
- New weapons science visualizations for the Thermonuclear Burn Initiative program
- Other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs
- PoP and GMV with modest updates to fix bugs and improve portability and maintainability

Preliminary planned activities in FY10:

- Continue to promote new discoveries in weapons science by advanced applications of visualization and data analysis in programs such as the Thermonuclear Burn Initiative
- Continue to document the results of our activities with classified papers and publications on weapons science topics jointly co-authored with X Division designers

WBS 1.5.4.6-SNL-001 Remote Petascale Data Analysis

The Remote Petascale Data Analysis project enables ASC customers to analyze, explore, and understand complex data results within CSSE's responsive design and analysis capability. A prominent theme of near-term work is support of V&V requirements. R&D in this project addresses a spectrum of needs, from tools promoting comparison and analysis of large numbers of runs to support for advanced analytic components needed by ASC simulations.

This project focuses on providing advanced customer-centered capabilities within an open source production framework (ParaView). Our scalable tools allow investigation of data on a variety of platforms—everything from a laptop to a cluster. This allows users to interact with their data, whether it fits on a PC or is located on a remote cluster in another state. In addition, our tools are beginning to deliver advanced analysis capabilities, in addition to traditional scientific visualization, that promote understanding of large data and investigation of ensembles of runs necessary for ASC's V&V goals.

In addition to providing these advanced capabilities, this project provides deployment and support services that enable ASC customers to carry out data analysis on both Zia and Trinity systems. This includes providing technical assistance for capability runs on the ACES platforms, as well as bridging the gap between advanced research and development and the users who must work with the entire end-to-end modeling and simulation environment.

Planned activities in FY09:

- Release initial advanced architecture support for analyses of complex data to promote investigation of ensembles of runs as well as single petascale-sized simulation results
- Provide advanced characterization tools for comparing large-scale simulated and experimental data

- Partner with customers in the SNL V&V program on advanced tool development for sensitivity analysis and UQ
- Provide ongoing technical and user support, including direct analysis support for capability class systems

Expected deliverables in FY09:

- Level 2 milestone to provide "Scalable Analysis Tools for Sensitivity Analysis and UQ"
- Preliminary implementation of Latency Tolerant Visualization capability, in support of FY10 GA of Zia platform
- Ongoing releases and installations of ParaView

Planned activities in FY10:

- Deliver advanced analysis capability in support of GA of Zia platform, including high-performance distance visualization, platform-specific optimization, and initial *in-situ* analysis capability (these areas align with ASC FY08 Petascale Data Analysis Milestone document)
- Deliver joint LANL/SNL Level 2 milestone, providing advanced hardware solutions for large data analysis in support of ACES
- Deliver initial multi-threaded and multi-core optimizations for scalable analysis tools
- Develop scalable methods of feature characterization and comparison tools in support of ACES

WBS 1.5.4.6-SNL-002 Visualization Deployment and Support

The Visualization Deployment and Support project deploys and supports software and infrastructure to deliver production data analysis and visualization capabilities. This project works closely with the Scalable Data Analysis project to provide a bridge from development to production; with many ASC projects to deliver its tools as part of an integrated modeling and simulation capability, including tri-lab; and with ASC platform and other infrastructure deployment activities to deliver a complete end-to-end system environment. This project also interfaces with applications and users to refine and enable use of ASC capabilities.

The project provides installation, testing, benchmarking, and end-user support for post-processing tools and utilities. Installation and support span the diversity of platforms in the distributed tri-lab environment. The project works with the other ASC labs to ensure that expected tools are deployed as needed to support the tri-lab computing environment. As to infrastructure, the project provides scalable data analysis systems that interactively analyze, visualize, and store output from ASC computers.

Planned activities in FY09:

- In partnership with the Scalable Data Analysis project, support delivery of ParaView/VTK-based tools that include enhancements for: petascale-sized simulation results; comparison of large scale simulated and experimental data; and sensitivity analysis and UQ
- Support infrastructure planning and evolutionary deployment toward petascale environments and, more specifically, the Zia platform environment

- Support the testing and performance analysis of data analysis software on multi-core general processing systems, for example, on TLCC
- Provide ongoing operational support of data analysis system infrastructure
- Provide ongoing user support, including direct visualization support for capability-class simulations

Expected deliverables in FY09:

- New release(s) of ParaView
- Level 2 milestone: "Scalable Analysis Tools for Sensitivity Analysis and UQ" (in partnership with Scalable Data Analysis project)

Preliminary planned activities in FY10:

- Deploy post-processing tools and associated local visualization/data-storage infrastructure needed to support SNL's use of capability computing, including Zia and Sequoia
- Provide ongoing delivery and support of up-to-date post-processing capabilities within an integrated responsive design and analysis capability

WBS 1.5.4.7: Common Computing Environment

The goal of the Common Computing Environment (CCE) product is to enable such an environment across the tri-labs that will initially be deployed on the TLCC systems. The scope of this product includes funded R&D projects to address gap areas identified by the tri-lab technical Working Groups.

The CCE working groups and projects focus on a common software stack to include, but not be limited to, operating system software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues of configuration management, licenses, WAN access, and multi-realm security, to name a few.

Common Computing Environment Deliverables for FY09

- Deliver tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"
- Deploy TOSS 1.2 (Based on RHEL 5.3)
- Implement consistent configuration management for TOSS, including evidence of a single code repository, use of this repository to manage releases of TOSS, and use of an appropriate release process including adequate testing
- Deploy tri-lab productivity on demand (Tripod) capabilities, including O|SS, Shared Work Space, and Gazebo Test and Analysis Suite
- Deploy tri-lab performance monitoring tools that provide consistent analysis and reporting of system usage and workload characterization across the tri-labs
- Implement production baseline of application monitoring tool with prototype of advanced implementation

WBS 1.5.4.7-TRI-001 Tripod Operating System Software

The TOSS is the tri-lab software stack to run across all newly procured Linux capacity clusters, initiating with TLCC platforms delivered in FY08. The goal of the TOSS project is to increase efficiencies in the ASC tri-lab community with respect to both the utility and the cost of the CCE. This project delivers a fully functional cluster operating system (kernel, Linux distribution, IB stack and related libraries, and resource manager) capable of running MPI jobs at scale on TLCC hardware. The system is to meet Tripod requirements for proving a common software environment on TLCC hardware across the tri-lab complex, now and into the future.

TOSS provides a complete product with full life-cycle support. Well-defined processes for release management, packaging, QA testing, configuration management and bug tracking are used to ensure that a production-quality software environment can be deployed across the tri-lab in a consistent and manageable fashion.

Planned activities in FY09:

- Provide ongoing TOSS software development and support
- Gather requirements and plan for TOSS 2.0 (to be based on RedHat RHEL 6)
- Support the tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Deployed TOSS 1.2 (Based on RHEL 5.3)
- Consistent configuration management for TOSS including evidence of a single code repository, use of this repository to manage releases of TOSS, and use of an appropriate release process including adequate testing
- Reasonable support processes including active and consistent issue tracking
- Appropriate coordination amongst the labs for overall TOSS support, such as support meetings happening on regular intervals

Preliminary planned activities in FY10:

- Provide ongoing TOSS software development and support
- Deploy TOSS 1.3 (Based on RHEL 5.4)
- Develop beta release of TOSS 2.0, contingent on RedHat RHEL 6 schedule

WBS 1.5.4.7-TRI-002 Open|SpeedShop

O|SS is targeted to be the main performance analysis tool set for the tri-lab ASC capacity machines. It is being developed jointly between the tri-lab partners and the Krell Institute. It provides many typical performance analysis steps in a single environment, including basic profiling in various forms, as well as MPI, I/O, and floating point exception tracing.

The O|SS efforts within Tripod include both maintenance and stability improvements to ensure the toolset's usability by our tri-lab user community as well as the design, implementation, and integration of new tools and scalability enhancements matching the demands of both tri-lab applications and TLCC hardware. The scalability enhancements will allow the project to investigate O|SS for ASC capability systems.

Planned activities in FY09:

- Create automated regression test and validation mechanism
- Increase scalability of tri-lab tool prototypes integrated with Tripod FY08 support
- Investigate O|SS for capability systems using the scalability enhancements
- Continue maintenance of O|SS on TLCC platforms
- Support the tri-lab Tripod Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Automatic regression testing on tri-lab capacity platforms
- Increased scalability for newly integrated tri-lab tools
- Prototype version of O|SS with limited functionality for capability systems

Preliminary planned activities in FY10:

- Provide further scalability enhancements to keep up with TLCC platforms
- Implement new memory tools focusing on non-uniform memory access aspects on TLCC systems
- Continue maintenance of O|SS on TLCC platforms
- Develop early production versions of O|SS for capability systems

WBS 1.5.4.7-TRI-003 Workload Characterization

The Workload Characterization project will develop a tri-lab common reporting interface for compute resource requirements for current and future use (with programmatic characterization of the work), and for platform usage data, tied to the programmatic characterization of the work.

Development and integration of tri-lab performance monitoring tool(s) may include: developing new functionality in SNL's HPC Estimations and Requirements Tool (HERT); modifying existing local laboratory tools; and integrating HERT with the Moab resource manager and, possibly, local laboratory tools and databases.

SNL's HERT web-based tool is viewed as the prototype for collecting and reporting, current and future requirements for compute resources, with programmatic characterization of the work. Additional development is needed in the areas of validation, test suites, and common reporting capability; as well as a more general mechanism that can interface to multiple existing databases at each of the tri-labs.

Resource Management/Moab development will be used to tie HERT estimates, with their respective workload characterization, to job requests and resulting platform usage data.

Planned activities in FY09:

- Develop new functionality and a common reporting interface based on SNL's web-based HERT tool; modify existing local laboratory tools, as required; and integrate HERT with local laboratory tools and databases
- Use Resource Management/Moab to tie HERT estimates, with their programmatic characterization of the work, to job requests and resulting platform usage data

- Support the tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Tri-lab performance monitoring tools that provide consistent analysis and reporting of system usage and workload characterization across the tri-labs

Preliminary planned activities in FY10:

- Assess usage and suggest enhancements in tools or documentation and training

WBS 1.5.4.7-TRI-004 Application Monitoring

The Application Monitoring project will develop tools that facilitate automated monitoring of production applications on ASC systems.

The tools should provide basic information about a user's job, to answer questions such as the following:

- Is the job making progress?
- How frequently is the job interrupted?
- What are causes and symptoms of interruptions to the job?
- Should the system intervene (for example, to kill or restart the job)?
- Should the system operators or the user be notified?
- How much time and storage is spent preparing for job restarts?

The project will develop a basic set of monitoring tools, along with their system and application interfaces. The result will be an extensible tool that can serve as a framework for future application monitoring functionality.

Planned activities in FY09:

- Develop and deploy baseline implementation on TLCC with intent to deploy on capability platforms
- Support the tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Production baseline implementation of application monitoring tool with prototype of advanced implementation

Preliminary planned activities in FY10:

None. This project is completed in FY09.

WBS 1.5.4.7-TRI-005 Shared Work Space

This project will deploy a collaborative on-line environment in which team areas can be created, with tools for communication, document reference, and other project centric tools to support planning and implementation. Ease of access from all of the tri-lab member locations is required.

Capabilities include ability to create team areas; posting of documents; documentation, plan, and reference material; code source repository access; task manager; access via current lab crypto-type cards; and Wiki, forums, email tracking and other tools.

Planned activities in FY09:

- Assess lab access policies and access abilities
- Choose collaborative environment and establish the environment on a server, accessible from all three labs using current authentication processes
- Conduct training and socialization at all three labs
- Assess usage and determine future needs
- Support the tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Production deployment of the collaborative, on-line, tri-lab shared environment with team areas and tools for communication, document reference, and project management

Preliminary planned activities in FY10:

- Support the collaborative, on-line, tri-lab shared environment with team areas and tools
- Assess one year of usage and determine desired enhancements

WBS 1.5.4.7-TRI-006 Gazebo Test and Analysis Suite

Gazebo is a collection of software components used to test, monitor, and analyze the health of a HPC system. With Gazebo, suites of system and application tests are run on an HPC system through either a web-based interface or from the system's master control node.

Test results are stored to a file system, and optionally to a database, and "normalized" so that a known set of timings and results establish a baseline for a healthy system. Through a set of tools, system analysts can monitor the health of the target running system and easily detect anomalous behavior.

Capabilities include results and coverage reporting tools (Command Line Interface only); database server and results schema; simple client-server communication protocol for network interaction; server daemon mythd (my test harness daemon); limited web client used for proof of concept; and acceptance test package (Command Line Interface).

This project will fully integrate CBENCH suite of tests into Gazebo. CBENCH is SNL's suite of test programs and scripts, which interrogate and report the status of individual hardware components comprising the cluster.

Planned activities in FY09:

- Provide updates or bug fixes as needed to ensure operability across tri-lab computing systems, specifically focus on configuration strategies to support inter-lab cluster variances.
- Define and create web-based interface(s) to fully support CBENCH test reporting requirements
- Interface output results with LANL monitoring project to begin producing base-line data for system anomaly correlation methodologies.
- Support the tri-lab Level 2 milestone "Deploy Tripod capabilities for capacity computing environment"

Expected deliverables in FY09:

- Production grade test harness and reporting tool
- Web-based interface and reporting tool(s) tailored for CBENCH and other tri-lab developed test suite requirements

Preliminary planned activities in FY10:

- Develop new or interact with other new technologies (system health monitoring tools and log analysis tools) to improve system root-cause analysis between failing or poor performing tests and the underlying system problem
- Position and deploy Gazebo as an open-source product to reap benefits of worldwide collaboration efforts

WBS 1.5.5: Facility Operations and User Support

This sub-program provides both necessary physical facility and operational support for reliable production computing and storage environments as well as a suite of user services for effective use of ASC tri-lab computing resources. The scope of the facility operations 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. FOUS 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. Industrial and academic collaborations are an important part of this sub-program.

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 FY09

- Deployment of the tri-lab common software environment on the FY08 TLCC systems
- Delivery of classified and unclassified capability and capacity computing cycles for the tri-lab, with unclassified time for the Academic Strategic Alliance Program (ASAP) and PSAAP partners
- Integration and deployment of common workload characterization software on LANL ASC platforms running the Moab resource management tools
- Development and delivery of a common report format for metrics on ASC platforms at LLNL, LANL, and SNL
- Roadrunner Phase 3 system integration into LANL network environment, including the execution of unclassified science runs on Roadrunner Phase 3 as part of system stabilization efforts
- Upgrade of DisCom encryptor to 10-Gig Ethernet unit
- Deployment of HPSS v7.1 in yellow and red networks
- Upgrade of the existing mechanical and electrical infrastructure in the SCC in preparation for the Zia capability system
- Site and prepare the infrastructure for Zia system

WBS 1.5.5.1-LLNL-001 System Administration and Operations

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.

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 a 24/7 operational monitoring capability for unclassified and classified computing environments consisting of large-scale computing platforms, infrastructure components, and networks. This project also covers self-maintenance of Linux clusters as well as spares inventory and control functions.

Planned activities in FY09:

- Prepare for the Sequoia ID (Dawn) system
- Support Purple, BlueGene/L, Peloton and TLCC capacity systems
- Move part or all of BlueGene/L from the classified to unclassified network
- Retire the Lilac SCF capacity cluster
- Retire the ALC capacity cluster, which will eliminate support of x86 production compute architectures at LLNL

Expected deliverables in FY09:

- Dawn system integration and support
- All or part of BlueGene/L moved to unclassified network
- Ongoing support for Purple, BlueGene/L, Peloton and TLCC capacity systems

Preliminary planned activities in FY10:

- Support Purple, BlueGene/L, Peloton and TLCC capacity systems
- Prepare for deployment of Sequoia

WBS 1.5.5.1-LLNL-002 Software and Hardware Maintenance, Licenses, and Contracts

The Software and Hardware Maintenance, Licenses, and Contracts project provides for vendor-provided hardware and software maintenance, support, licenses, and development contracts. For laboratory-maintained systems, the project provides hardware maintenance capabilities including component inventory and replacement.

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 FY09:

- Track and place contracts and licenses needed for system operations and vendor support
- Ongoing hardware self-maintenance of BlueGene/L, parallel global file systems, Peloton, and TLCC capacity systems

- Self-maintenance of Dawn

Expected deliverables in FY09:

- Contracts and licenses needed for system operations and vendor support put in place
- BlueGene/L, parallel global file systems, Peloton and TLCC capacity systems properly maintained
- Self-maintenance of Dawn
- Dawn spares inventory completed

Preliminary planned activities in FY10:

- Put in place contracts and licenses needed for system operations and vendor support
- Properly maintain BlueGene/L, Dawn, parallel global file systems, Peloton and TLCC capacity systems

WBS 1.5.5.1-LLNL-003 Computing Environment Security and Infrastructure

The Computing Environment Security and Infrastructure project provides for the development, enhancement, integration, and ongoing support of the core security infrastructure services and cyber-security environment, including security-enabled software components and interfaces necessary for the efficient, effective, and secure use of large-scale ASC platforms by local and remote customers. This activity also involves developing, porting, and testing the security middleware software stack that enables and enforces centralized authentication, access control, and data sharing for users of these platforms.

The project deploys and supports the centralized authentication, authorization, and security registry services for ASC platforms in the unclassified and classified networks. Integral to this is the ongoing development and integration of the cyber-security infrastructure and supporting services, including but not limited to, backups, collaboration and productivity tools, monitoring and cyber-security management tools, and security middleware software stack for ASC platforms. Project personnel participate in the design, development, integration, and management of this robust infrastructure environment to support large-scale ASC platforms.

Planned activities in FY09

- Continue upgrades to core security infrastructure components, thus enabling enhanced security capabilities and mechanisms to be offered and utilized
- Integrate security middleware software stack on Dawn ID system and new TLCC cluster platforms
- Develop and integrate additional electronic workflow and provisioning functionality that leverages the deployed identity management solution
- Refine the existing operations methodology of the cyber-security environment in the unclassified and classified computing environment
- Perform ongoing security-related activities in support of the secure use and secure management of ASC platforms and associated infrastructure

Expected deliverables in FY09

- Deployment of lightweight directory access protocol/one-time password replacement for one-time password service and complete transition
- MIT Kerberos authentication service and lightweight directory access protocol directory services upgraded to latest compatible releases
- Deployment of identity management phase II features, including privileged account management, bulk operations and electronic provisioning for the classified environment
- Security architecture blueprint document written

Preliminary planned activities in FY10

- Investigate and evaluate approaches and technologies identified in the security architecture blueprint
- Continue to incorporate external data sources and evaluate process improvements for a comprehensive identity management system

WBS 1.5.5.1-LLNL-004 Facilities Infrastructure and Power

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, funding needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff is part of this project.

Planned activities in FY09:

- Continue to track the progress of the institution project elements for increased redundancy and reliability of the laboratory electrical distribution system that will support the 15-megawatt electrical power expansion for B-453 for 2011 completion
- Maintain and support equipment in existing computational and staff facilities
- Continue analysis of future modifications and/or expansion of facilities that will be needed by future ASC systems
- Continue to utilize self-benchmarking tool for all of the computer rooms in B-453 (TSF), B-451, B-439, B-115, and B-117 created by DOE Office of Science and LBL (http://hightech.lbl.gov/datacenters.html) to routinely identify prospective energy savings initiatives as computer rooms change
- Complete baseline computational fluid dynamics model for all ASC systems to profile the airflow required to cool the machines adequately

Expected deliverables in FY09:

- Site preparations for the Sequoia ID system and Hyperion systems
- Electrical procurement of equipment and construction of the first phase of the B-453 15-megawatt expansion
- Final cooling tower cell installation at U-454
- Final 1200 ton chiller installation at B-453

Preliminary planned activities in FY10:

- Continue to track the progress of the institution project elements for increased redundancy and reliability of the laboratory electrical distribution system that will support the 15-megawatt electrical power expansion for B-453 for 2011 completion
- Complete the first phase of construction of the B-453 15-megawatt power expansion
- Start the second phase of construction of the B-453 15-megawatt power expansion
- Continue to maintain and support the equipment in existing computational and staff facilities
- Continue analysis of future modifications and/or expansion of facilities that will be needed by future ASC systems

WBS 1.5.5.1-LLNL-005 Classified and Unclassified Facility Networks

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 include 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.

Planned activities in FY09:

- Provide network architecture design and testing to support the Sequoia ID system
- Enhance core network 10-gigabit Ethernet density to support follow on TLCC platforms and HPSS networks
- Evaluate the use of alternative technologies for use as network interconnect for future systems
- Evaluate alternatives for moving Open LabNet (OLN) connection to 10 Gigabit
- Provide ongoing local network performance, reliability, security, and operational support

Expected deliverables in FY09:

- Network connection deployment for the Sequoia ID system
- 10 Gigabit connectivity to the HPSS network deployment

Preliminary planned activities in FY10:

- Provide network architecture design and testing to support the Sequoia platform
- Evaluate use of 40 Gigabit connectivity in the backbone
- Provide ongoing local network performance, reliability, security, and operational support

WBS 1.5.5.1-LLNL-006 Wide-Area Classified Networks

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.

Capabilities include ongoing discussions with the tri-lab community, which 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 WANs 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.

Planned activities in FY09:

- Track development and availability of high-speed encryption units for use in the classified WAN
- Provide operational support of the tri-lab WAN connections

Expected deliverables in FY09:

- 10 Gigabit encryption in the DISCOM WAN deployment

Preliminary planned activities in FY10:

- Track development and availability of 40 Gigabit encryption and beyond for WAN use between petaFLOPS platforms
- Provide operational support of the tri-lab WAN connections

WBS 1.5.5.1-LANL-001 High-Performance Computing Operations Requirements Planning

The HPC Operations Requirements Planning project covers the planning activities for computing operations, collection, and statistical evaluation of user requirements for computing resources, development of new metrics, and data collection.

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

Planned activities in FY09:

- Coordinate with other nuclear complex laboratories to design, develop, and deploy common workload characterization software. The software will depend on a Moab resource management installation and customer input into a workload characterization database. The system will be integrated with the existing HPC account management at LANL.

Expected deliverables in FY09:

- Integration and deployment of common workload characterization software on LANL ASC platforms running the Moab resource management tools.
- Development and delivery of a common report format for these platforms.

Preliminary planned activities in FY10:

- Continue to contribute to improvements in the quality and reliability of computing requirements collection and analysis, both at LANL and in the tri-lab arena.

WBS 1.5.5.1-LANL-002 Roadrunner Phase 3 Initial Deployment

The scope of the Roadrunner Phase 3 Initial Deployment project is to take delivery and start the deployment of the Roadrunner Phase 3 system. This includes completing the acceptance tests, system integration into the LANL network, system stabilization, and transition into the classified network. This section maps to the Roadrunner Project Element "3.3 Final System Integration."

The primary capabilities are acceptance and diagnostic testing, system stabilization, system integration into the yellow network, and transition to secure network.

Planned activities in FY09:

- Complete acceptance testing
- Evaluate delivered system with diagnostic testing
- Conduct system tests on the overall compute system
- Plan for integrating the system into secure computing environment

Expected deliverables in FY09:

- Completion of Roadrunner Phase 3 acceptance tests
- Roadrunner Phase 3 system integration into LANL network environment
- Completion of science runs on Roadrunner Phase 3 as part of system stabilization efforts
- Completing of security and test plans
- Transition Roadrunner Phase 3 into classified network

Preliminary planned activities in FY10:

- Transition system to production status
- Develop integrated support structure with vendor for operational issues

WBS 1.5.5.1-LANL-003 Ongoing Network Operations

The Ongoing Network Operations project provides ongoing network operations that support ASC computing in the classified and unclassified networks. This includes directly attached networks to HPC systems (machine area network), network backbones, user LANs, and the high-end DisCom WAN connecting the tri-labs.

Core capabilities include designing, developing, deploying, and supporting classified and unclassified network hardware and services to support ASC computational systems and infrastructure.

Planned activities in FY09:

- Improve data delivery for scientific visualization
- Operate and maintain network services
- Support InfiniBand interconnect fabrics

- Manage the high-performance network backbone
- Refine network to improve performance and availability of backbone and services
- Support new NNSA Enterprise Secure Network

Expected deliverables in FY09:

- New technology deployment to increase network utilization, bandwidth, and reliability
- Internet protocol encryptor upgrade
- Ongoing network support for all ASC systems

Preliminary planned activities in FY10:

- Continue to operate and maintain LAN/Metropolitan Area Network/WAN infrastructure
- Additional integration with Enterprise Secure Net (ESN) network
- Plan I/O infrastructure to support the Zia platform

WBS 1.5.5.1-LANL-004 Network Infrastructure Integration

The Network Infrastructure Integration project includes all services for networks operated by LANL for the purpose of providing a HPC networking environment for weapons designers, developers, and engineers.

Core capabilities include designing, procuring, prototyping, testing, and installing network hardware and software to meet ASC bandwidth and performance requirements.

Planned activities in FY09:

- Design, plan, and start integration of the network infrastructure to support the new Zia capability system and its integration into the Integrated Computing Network infrastructure
- Design and integrate the network infrastructure to support the additional TLCC clusters
- Extend our GPFS networks, parallel scalable back bone concept, to support new HPC clusters and I/O links in both the Yellow and Red networks
- Upgrade network infrastructure, where necessary, to increase bandwidth and availability to HPSS, visualization platforms, and customer workstations
- Provide support for system interconnects monitoring, reliability and performance management

Expected deliverables in FY09:

- Network enhancements for new TLCCs
- Network optimization to improve performance to Roadrunner Phase 3
- Design completion of I/O Network Integration for Zia
- DisCom encryptor upgrade to 10-Gig Ethernet unit

Preliminary planned activities in FY10:

- Complete the integration of the Zia system into the network infrastructure in preparation for full-scale operation

WBS 1.5.5.1-LANL-005 Ongoing Systems Operations

The Ongoing Systems Operations project includes all services for systems operated by LANL for the purpose of providing an HPC production computing environment for weapons designers, developers, and engineers. The project works with users to troubleshoot problems experienced while running their applications and helps users transition from old to new computing platforms.

The capabilities provided include system configuration, system and user security, resource management, system administration, system operation, monitoring, and hardware maintenance.

Planned activities in FY09:

- Support system by conducting ongoing and daily system administration with continuous monitoring of production systems and infrastructure servers
- Ensure workload is carried out by proper configuration of queues and scheduling policies; daily monitoring and problem resolution of use problems.
- 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
- Build on reliability, availability, and serviceability (RAS) metrics capabilities delivered in FY08
- Deploy monitoring infrastructure to Roadrunner full system and continue infrastructure improvements for Roadrunner base and TLCC platforms
- Provide around-the-clock operations and monitoring of the scientific computing resources, including an increased level of system hardware self-maintenance for various computing and storage systems
- Continue to study the feasibility of reduced staff ("lights out") and of a "skewed" work week for computer operation's staff
- Implement additional formal and informal customer satisfaction metrics and measurement techniques
- Expand hardware self-maintenance to additional systems and components as practical
- Ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Deliver a reliable monitoring system to the computer operations staff for monitoring scientific computing machine status as well as early detection of hardware problems
- Train operations staff in the hardware maintenance of Roadrunner and TLCC
- Provide hardware support for installation and integration of additional scientific computing platforms, including Roadrunner and TLCC

Expected deliverables in FY09:

- Transition of Roadrunner Phase 3 system into full production
- Installation and deployment of TLCC SUs for production use including Tripod software and TOSS stack
- HPSS v7.1 deployment in yellow and red networks
- Monitoring infrastructure deployment to Roadrunner Phase 3 platform
- Preparation for and initial installation of Zia

Preliminary planned activities in FY10:

- Provide support for Roadrunner Phase 3 Weapons Science calculations and for advanced architecture studies with integrated weapons applications
- Complete installation of Zia and transition to GA
- Deploy a more extensive use of Tripod tools and capability on the TLCC system

WBS 1.5.5.1-LANL-008 Ongoing Facilities

The Ongoing Facilities project is responsible for the engineering, design, operation, and maintenance of the electrical, mechanical, cooling, and other computing infrastructure in support of the ASC program. Electrical costs for capability and capacity platforms are paid for by funding in this project. A major activity in FY09 will be the power and cooling infrastructure equipment upgrades to the SCC in support of the Zia machine.

The primary capability of this project is supporting LANL's ASC computing facilities by providing the necessary engineering, design, and maintenance activities that enable ASC HPC platforms and associated infrastructure to operate at full capacity.

Planned activities in FY09:

- Upgrade existing mechanical and electrical equipment in the SCC in preparation for the Zia capability system
- Create a five-year plan addressing site prep requirements, including power, cooling, and space for current and future HPC installations
- Continue operations and maintenance of electrical and mechanical systems for ASC computing facilities.
- Provide facility support for decommissioning of retired systems
- Increase power use effectiveness and data center efficiency by implementing engineering strategies in an effort to optimize energy efficiencies and cost savings in our computing facilities
- Develop a long-term project planning tool for future of ASC computing projections with input from power and cooling subject matter experts and other supercomputing expertise

Expected deliverables in FY09:

- Significant progress towards the completion of the infrastructure equipment upgrade in the SCC in support of the Zia machine
- Site preparation for the Zia system in the SCC is complete
- Five-year facilities plan

Preliminary planned activities in FY10:

- Complete final infrastructure equipment upgrade in the SCC in support of the Zia machine
- Begin site preparations for installation of future computing platforms
- Continue operations of the SCC, Laboratory Data Communications Center (LDCC), and Central Computing Facility (CCF) facilities and computer rooms
- Plan infrastructure improvements to keep power and cooling capabilities commensurate with new supercomputer power and cooling requirements
- Make enhancements to electricity and cooling to support expected increase in computing capacity in the LDCC along with new system cooling requirements

WBS 1.5.5.1-SNL-001 Production Computing Services

The Production Computing Services project's goals are to operate and maintain all production platforms and associated support systems, and operate ASC capability and capacity platforms, data services and visualization systems, long-term hierarchical storage services, high-performance network systems, tri-lab compatible cyber authentication and authorization systems, and monitoring and reporting services. This project will support tri-lab capability platform resource allocations and coordinate with tri-lab peers in establishing priority scheduling if required. This project coordinates the integration and deployment of TLCC capacity systems into SNL's production computing environment, in collaboration with WBS 1.5.4.7 Common Computing Environment. Support of Tripod common service and environment decisions and configuration management activities will also be provided.

Planned activities in FY09:

- Operate production systems in support of nuclear weapons mission requirements—capability, capacity, data services and visualization systems, long-term hierarchical storage
- Bring to production state three TLCC07 capacity systems: Unity, Whitney, Glory
- Retire obsolete SNL Linux clusters by replacing them with TLCC clusters
- Transition data from end-of-life tape systems to StorageTek SL8500s
- Decommission SGI visualization servers and clusters based in CA

Expected deliverables in FY09:

- Three TLCC capacity production systems running Tripod/TOSS common software stack
- Reliable monitoring system for the networking infrastructure connecting ASC tri-lab classified computing environments, integrated with the new ESN infrastructure
- Network switching infrastructure enabling time sharing of Red Storm between multiple TS environments
- Upgraded Kerberos Services enabling Smartcard authentication

Preliminary planned activities in FY10:

- Research and prototype modifications to Kerberos Service to implement policy-based authorization

WBS 1.5.5.1-SNL-002 Facilities and Infrastructure

The Facilities and Infrastructure project will manage all computing facilities and support infrastructure, and provide funds for physical security and utilities (power, cooling, and space) expenses. It will also plan and coordinate facilities construction or expansion; design and procure power and cooling equipment as required for production platforms; supply physical security control for computing facilities and classified media; and provide 24 hour per day support. Finally, this project will provide required support for conduct secure operations.

Planned activities in FY09:

- Assess and plan acquisitions for incremental power, cooling and facility modifications to support future-year TLCC installations
- Complete business plan for B-725 expansion
- Complete transition of B-725 to support enhanced compartmented security
- Manage retirement and destruction of classified disk and tape media

Expected deliverables in FY09:

- Fully compliant facility for HPC needs of DOE/NNSA and other government agencies supporting the Field Intelligence Element at SNL
- Reduced power consumption for air handling equipment through use of new DC powered York units procured in FY08 and FY09 (expected reduction of 30 percent of air handler power)

Preliminary planned activities in FY10:

- Develop alternative energy options for data centers and existing computing facilities
- Invest in data centers that can accommodate the growing power/cooling needs of newer HPC technology
- Deploy networking infrastructure on classified to mirror infrastructure on restricted environment (Force10 switch)

WBS 1.5.5.1-SNL-003 Tri-Lab System Integration and Support

The Tri-Lab System Integration and Support project manages projects relating to tri-lab production networking services and related infrastructure. SNL provides coordination, operational support and oversight to develop and operate the ASC WAN and manages the communication link contracts. SNL leads the integration of new encryptor technology into the WAN by evaluating early engineering samples and organizing tri-lab wide functional testing prior to deployment.

Traffic engineering and modeling systems, as well as a dedicated test laboratory based WAN development environment, are used to improve network efficiency and utilization. Monitoring and management systems are used to analyze network performance and validate vendor availability data to ensure proper credits are applied to the communication link contracts. The project oversees the Qwest communication link contract and monitors ESNET connectivity. It provides system-level analyst support for cross-site production services related to data transfer, distance computing, and access methods and services; and coordinates production requests for tri-lab resources.

Planned activities in FY09:

- Support customer support/production operations and tri-lab customers

- Support platform monitoring and notification for cyber enterprise management
- Operate tri-lab ASC WAN and manage the Qwest communication link contract
- Test data movement tools and services between LANL and SNL for Zia
- Improve monitoring capabilities for data movement tools and services

Expected deliverables in FY09:

- ASC WAN tri-lab capability satisfying FY08 Office of Management and Budget mandate (increased bandwidth from 10GEthernet encryptors)
- Emergency tri-lab ESnet ISP access via ASC WAN
- 10-Gigabit network infrastructure service to the user desktop environment
- SecureNet conversion to Enterprise Secure Network
- Automated account process integrating SARAPE with WebCARS and providing NWC wide support (delayed from FY08)

Preliminary planned activities in FY10:

- Provide highly integrated monitoring of local and remote production platforms
- Negotiate new contract for WAN services
- Support data movement tools and services for Zia

WBS 1.5.5.1-Y12-001 Applications in Support of Manufacturing Production and Connectivity

The Applications in Support of Manufacturing Production and Connectivity project supports the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. The project includes support for connecting to ASC computing resources and job submission, execution, and visualization. The project also supports the transition of the Y-12 compute cluster to the classified production environment and provides the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources. Development and implementation of software to support the solution of manufacturing problems is also supported by the project. Visualization techniques that can be utilized in the Y-12 network and computing infrastructure will be evaluated and implemented. Finally, participation in Nuclear Weapons Complex ASC-related activities are covered.

The function of this project is to support the utilization of ASC codes, computing resources, and techniques to solve production manufacturing problems through modeling and simulation. The Y-12 compute cluster will be transitioned to the classified environment, and tools will be implemented to support visualization and software development and implementation.

Planned activities in FY09:

- Utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Refine plan to integrate Y-12 compute resource with ASC environment
- Participate in Nuclear Weapons Complex ASC activities
- Promote follow-on activities from Computational Manufacturing conference held at Y-12 in FY08

- Transition Y-12 cluster resource to classified production environment
- Support application deployment on classified cluster to meet Y-12 production requirements

Expected deliverables in FY09:

- A report describing research results on multiscale methods for computed tomography
- Modified computed tomography codes to accommodate large projection datasets
- Demonstrated material modeling code on cluster
- Demonstrated manufacturing optimization code on cluster

Preliminary planned activities in FY10:

- Utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Participate in Nuclear Weapons Complex ASC activities
- Investigate the use of uncertainty quantification techniques in tomographic reconstruction of stockpile components

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 FY09:

- New training classes and documentation developed for TLCC clusters with TOSS stack
- New training classes for Roadrunner computing with focus on efficient hybrid computing
- Initial roll-out of training and documentation Web sites
- Redesigned accounts system in production use
- Reliable and responsive service to users in the ASC tri-lab computing environments
- Training for system administrators and users via Web-based short courses
- SNL corporate-managed training and account synchronization
- Online support on classified network
- Customer support via tri-lab service level agreement (SLA) for transferring HPC support issues among labs

WBS 1.5.5.2-LLNL-001 Hotlines and System Support

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. Services are provided to users from external sites including LANL, SNL, and the ASC Alliance sites, as well as the LLNL users.

Planned activities in FY09:

- Transition user accounts processing to new identity/account management tool
- Deliver TLCC training to ASC and Alliance partners
- Develop Sequoia ID system documentation and training materials
- Upgrade or replace the existing trouble ticket system
- Provide ongoing support services for hotline operations, documentation, and training

Expected deliverables in FY09:

- Sequoia ID system documentation and training materials
- Trouble ticket system upgrade or replacement

Preliminary planned activities in FY10:

- Deliver Sequoia ID system training to ASC and Alliance partners
- Investigate new tools for delivering system status information to users
- Provide ongoing support services for hotline operations, documentation and training

WBS 1.5.5.2-LANL-001 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support

The Integrated Computing Network Consulting, Training, Documentation and External Computing Support project is responsible for direct customer service for local and remote users of ASC/LANL resources, the development and delivery of documentation and training materials for ASC/LANL resources, usage statistics, and an administrative interface for ASC tri-lab and Alliance users, and other external ASC/HPC users.

The primary capabilities consist of user support services, operational metrics for an HPC environment on, for example, usage and availability, Web page development to present this information to system personnel and users, and the development of user documentation and training.

Planned activities in FY09:

- Perform ongoing user support for users of ASC/LANL computing resources
- Expand online documentation in both breadth and depth
- Deploy an improved documentation and training set of Web sites, allowing for ASC user community interaction within the site
- Deploy a redesigned accounts system

- Refine and develop ongoing customer satisfaction surveys

Expected deliverables in FY09:

- New training classes and documentation developed for TLCC clusters with TOSS stack
- New training classes for Roadrunner computing with focus on efficient hybrid computing
- Initial roll-out of training and documentation Web sites
- Redesigned accounts system in production use

Preliminary planned activities in FY10:

- Perform ongoing user support for users of ASC/LANL computing resources
- Continue to improve the quality of support for productive use of computing resources
- Explore alternative communication methods with users and automated monitoring processes to determine where such methods may be applicable to providing responsive support within the ASC program

WBS 1.5.5.2-SNL-001 User Environment and Application Support

The User Environment and Application Support teams facilitate computing on ASC tri-lab platforms, as well as SNL computing systems. User support activities are focused on improving the productivity of the entire user community, local or remote, in utilizing the ASC HPC resources. The HPC user environment is complex and with the introduction of new advanced architecture systems is becoming much more diverse. Different compute platforms, different compiler environments, different file systems, different operating systems on the resources, and constantly modified applications codes create a multi-dimensional problem space which requires experienced, dedicated, and innovative user support personnel to identify and correct faults discovered within the environment. These tasks are best addressed by support personnel who are trained in HPC rather than the application code developers who focus more on application code improvements.

The User Environment and Application Support project provides information, tools, training, and direct user support for the ASC scientific computing environment. Resources include a knowledge management and retrieval system (collaborative learning, information, and knowledge) email and Web-enabled "self-support" tool. User support is provided via prime-time telephone and email support. Tri-lab support includes assisting tri-lab customers with problems at SNL; assisting SNL customers with computing at remote locations; management of SNL computing resources via the CCC process for Purple and the SNL Platform Oversight Committee process; and representing SNL needs to the expedited priority run process.

The application support team works with a breadth of ASC applications and system environments to develop and apply expertise that enables efficient and effective use of ASC's precious computing resources. The team provides porting assistance to new ASC architectures; benchmarking; and various forms of application performance analysis, including modeling, scaling studies, and optimization.

Planned activities in FY09:

- Improve reliability and responsiveness to ASC tri-lab and SNL users' requests for assistance; aim is to improve analyst efficiency, effective use of ASC resources, and accuracy of modeling and simulation result
- Design an operational support model for Zia (in partnership with LANL) that services SNL users and applications; aim is to minimize the disruption to code team development practices and analyst work effort introduced by a remote capability computing system
- Provide TLCC documentation and training; object is to improve the effectiveness of the principle capacity computing systems in supporting ASC projects
- Model SNL applications to evaluate and improve performance for future advanced architectures
- Optimize SNL application performance for Red Storm and TLCC quad-core nodes to ensure high return on investment
- Support SNL applications for large capability runs on Red Storm and Purple; full system simulation runs often expose weaknesses within the operating system, the file system, the user environment (for example, compilers, libraries, debuggers, profilers), and the applications which limit the scale of a simulation
- Manage SNL computing resources via tri-lab expedited priority run meetings, CCC process, and Sandia Platform Oversight Committee process at SNL; ensure most efficient and equitable use of ASC resources

Expected deliverables in FY09:

- Reliable and responsive service to users in the ASC tri-lab computing environments
- Training for system administrators and users via Web-based short courses
- SNL corporate managed training and account synchronization
- Online support on classified network
- Customer support via tri-lab service-level agreement for transferring HPC support issues among labs

Preliminary planned activities in FY10:

- Implement the joint LANL-SNL operational support model for Zia
- Provide porting assistance to new ASC platforms/architectures (including Zia and Sequoia); benchmarking; and application performance analysis, including scaling studies, and optimization
- Manage SNL computing resources via tri-lab expedited priority run meetings, CCC process, and Sandia Platform Oversight Committee process

WBS 1.5.5.3: Collaborations

This level 4 product provides collaboration with external agencies on specific HPC 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 FY09

- Support for Supercomputing conference
- Host the Predictive Science Panel
- Multi-media and written communications for NA-114; conceptualize, develop, provide, and manage
- Support for cooperative work with the SC09 ASC Research Exhibit

WBS 1.5.5.3-LLNL-001 Program Support

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

Planned activities in FY09:

- Manage existing tri-lab contacts and negotiating and executing new contracts
- Support the annual Supercomputing conference, Predictive Science Panel meetings, and other meetings and workshops
- Provide support to the ASC Federal program-management office

Expected deliverables in FY09:

- Continued management of existing Purple and BlueGene/L contracts as well as the research contract with IBM for BlueGene/P and BlueGene/Q
- Closeout of the original ASAP contracts
- Execution and management of common computing environment and TLCC-related contracts
- Negotiation and execution of Sequoia contract

Preliminary planned activities in FY10:

- Continue FY09 activities

WBS 1.5.5.3-LLNL-002 Scientific Collaborations

LLNL scientists, together with existing codes and resources at LLNL, have unique capabilities to address scientific challenges of interest to other parts of DOE and other government agencies. This project provides support for collaborations with the NNSA Office of Nuclear Nonproliferation, with the Office of Science for multi-institution SciDAC projects, and with the Defense Threat Reduction Agency (DTRA) Nuclear Weapons Effects Division.

Planned activities in FY09:

- Participate in a joint NNSA/Office of Science SciDAC project led by the University of Southern California to develop a hierarchical petascale simulation framework addressing stress, corrosion, and cracking in metals and alloys from first principles, with specific emphasis on the hybrid coupling of quantum simulations and quantum-based atomistic simulations to develop optimized potentials for SCC applications
- Participate in a joint NNSA/Office of Science SciDAC project led by Stanford to develop improved numerical methods for flows involving shocks, turbulence and

strong density gradients with special emphasis on ensuring the new methods scale to hundreds of thousands of processors

- Collaborate with DTRA Nuclear Weapons Effects Division on electromagnetic pulse effects modeling and simulation; LLNL codes will be installed at the DTRA Nuclear Weapons Effects Division Computational Multiphysics laboratory and investigated for coupling with existing DTRA electromagnetic pulse codes
- Collaborate with the Russian Federation Institutes

Expected deliverables in FY09:

- A plan for a suite of modern, integrated simulation codes for electromagnetic pulse generation, propagation, and system response

Preliminary planned activities in FY10:

- Continue FY09 activities

WBS 1.5.5.3-LANL-001 Program Support

Through the Program Support project, LANL provides support to the national program, both by providing resources and expertise to the Federal program office and by participating in coordination and integration activities for the tri-lab program.

Planned activities in FY09:

- Alternate with Livermore in hosting the Predictive Science Panel; results will be incorporated into program plans and initiatives
- Provide consultant support to the Federal program management efforts to foster collaborations and build support within the predictive science community
- Support for the PSAAP

Expected deliverables in FY09:

- Support for SC08 conference
- Host the Predictive Science Panel

Preliminary planned activities in FY10:

- Host the Predictive Science Panel
- Lead the organization of the ASC tri-lab booth at the SC09 conference

WBS 1.5.5.3-SNL-001 One Program/Three Labs

One Program/Three Labs funds critical coordination and integration activities essential to the success of ASC. These are divided into two distinct parts: 1) provide ASC multi-level communications per existent communications plan and by special request, and 2) SNL outreach to the DoD laboratories and programs.

Planned activities in FY09:

- Continue development of the SNL and HQ ASC Web sites
- Continue production of high-quality communications materials for HQ and the broader HPC community
- Perform special assignments for Complex Transformation, such as the S&T roadmap (as required) and special events (expositions)

- Support the ASC executive committee; support for quarterly meetings of the ASC executive committee; and management of the SAIC contract to provide various administration support to HQ
- Support multi-media and written communications for NA-121.2; conceptualize, develop, provide, and manage communications

Expected deliverables in FY09:

- Multi-media and written communications for NA-121.2
- Support for cooperative work with the SC09 ASC Research Exhibit

Preliminary planned activities in FY10:

- Continue development of the SNL and HQ ASC Web sites
- Continue production of high-quality communications materials for HQ and the broader HPC community
- Perform special assignments for Complex Transformation such as the S&T roadmap (as required) and special events (expositions)
- Support for the ASC executive committee; support for quarterly meetings of the ASC executive committee; support for the annual PI meetings that expose attendees to technical and programmatic efforts at the three laboratories; and management of the SAIC contract to provide various administration support to HQ

V. ASC Level 1 and 2 Milestones

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

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)	ASC Category
2	Develop, implement, and validate a suite of physics-based models and high-fidelity databases in support of Full Operational Capability in DTRA's National Technical Nuclear Forensics program.	1	FY09	Q4	ASC	HQ, LLNL, LANL	C11, C1, C4, NA-22, DTRA
1	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of the initial conditions for secondary performance.	1	FY10	Q4	ASC	HQ, LLNL, LANL	C11, C4
5	Assessment of weapon surety status (nuclear safety and physical security) in off-normal transportation scenarios.	1	FY12	TBD	ASC	HQ, SNL	
3	Baseline demonstration of UQ aggregation methodology for full-system weapon performance prediction	1	FY12	Q4	ASC	HQ, LLNL, LANL, SNL	C11, C1, C4, DSW
6	Demonstrate predictive capability for weapon system response to short-pulsed neutrons in hostile radiation environment.	1	FY13	TBD	ASC	HQ SNL	
7	Full-system safety assessment of damaged weapon immersed in fuel fire for transportation accident scenario.	1	FY14	TBD	ASC	HQ, SNL	
4	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of the initial conditions for primary boost.	1	TBD	TBD	ASC	HQ, LLNL, LANL	C11, C1, C2

Table V-2. Quick Look: Level 2 Milestone Dependencies for FY09[5]

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	Deliver enhanced above-ground experiments (AGEX) capability	2	09	Sep-09	IC	LLNL
TBD	Improve physics models, code usability, and nuclear databases in support of the full operational capability needed by the National Technical Nuclear Forensics program by Dec. 31, 2009	2	09	Sep-09	IC	LLNL
TBD	Improve HED AGEX modeling capability in LLNL's ASC nuclear design code system for validation of models required for the FY20 L1 milestone for predictive secondary performance simulation capability	2	09	Jun-09	IC	LLNL
TBD	Establish initial high-fidelity full-device baseline capability	2	09	Dec-08	IC	LLNL
TBD	New global EOS data library delivered for QMU, V&V, and other applications	2	09	Sep-09	PEM	LLNL
TBD	High-resolution exploration of two key UGT events	2	09	Sep-09	PEM	LLNL
TBD	Advanced high explosive detonation model to support enhanced predictivity in 3D simulations	2	09	Sep-09	PEM	LLNL
TBD	Top 20 uncertain parameters in a primary device calculation	2	09	Jun-09	V&V	LLNL
TBD	First Demonstration of Secondary Computational Assessment Metric Project (SCAMP) Across Several UGT's using the Computational UQ Pipeline	2	09	Sep-09	V&V	LLNL
TBD	TLCC clusters in production at LLNL	2	09	Dec-08	CSSE, FOUS	LLNL
TBD	Deploy Dawn ID machine for initial science runs	2	09	Sep-09	CSSE, FOUS	LLNL
TBD	Usability enhancements to support Directed Stockpile Work (DSW)	2	09	Mar-09	IC	LANL
TBD	2D high-fidelity burn capabilities to support Directed Stockpile Work (DSW)	2	09	Sep-09	IC	LANL
TBD	Release of a Crestone project code primarily focused on improving post-processing capability	2	09	Mar-09	IC	LANL
TBD	Release of the Crestone project codes primarily focused on optimizing Implicit Monte Carlo on the Roadrunner hybrid cell processor platform	2	09	Sep-09	IC	LANL
TBD	Release of merged MCNP5 and MCNPX code	2	09	Sep-09	IC	LANL
TBD	Roadrunner mission science simulations: ejecta and TN Burn	2	09	Sep-09	PEM	LANL
TBD	Attribution – Deliver ASC capabilities for nuclear forensics	2	09	Sep-09	PEM	LANL
TBD	Plutonium fission – models for neutron spectra and UQ assessment	2	09	Sep-09	PEM	LANL
TBD	Assessment of specific ASC code capabilities for Thermonuclear Applications	2	09	Jun-09	V&V	LANL

[5] Factors such as FY09 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)
TBD	Initial implementation of a LANL Boost Validation Suite (BVS)	2	09	Sep-09	V&V	LANL
TBD	Calculation verification of a full-system mechanical response simulation of a hostile blast environment	2	09	Sep-09	V&V	LANL
TBD	Prototype full-system uncertainty quantification using ASC code capabilities	2	09	Sep-09	V&V	LANL
TBD	Roadrunner system integration readiness	2	09	Jun-09	CSSE	LANL
TBD	Performance assessment of major tri-lab ASC platforms	2	09	Sep-09	CSSE	LANL
TBD	HPSS 7.1 deployment	2	09	Sep-09	FOUS	LANL
TBD	Completion of Phase 1 consolidation of thermal/fluid capabilities and consolidation of solid mechanics capabilities in SIERRA Mechanics	2	FY09	Sep-09	IC	SNL
TBD	Initial coupled thermal/fluid capability in SIERRA Mechanics to compute ablation and thermal response of thermal protection systems on a Reentry Body/Reentry Vehicle (RB/RV)	2	09	Sep-09	IC	SNL
TBD	Improved energy dissipation models for predictive mechanical response	2	09	Sep-09	IC	SNL
TBD	Initial demonstration of consolidated DART Workbench	2	09	Sep-09	IC	SNL
TBD	Improved pressure fluctuation model for turbulent flow	2	09	Sep-09	P&EM	SNL
TBD	Develop and implement in Aleph coupled electrode/plasma models for simplified arc initiation	2	09	Sep-09	P&EM	SNL
TBD	Algorithms for smart, adaptive uncertainty assessments	2	09	Sep-09	V&V	SNL
TBD	Uncertainty analysis for the prediction of x-ray dose-rate environment in a weapon system	2	09	Sep-09	V&V	SNL
TBD	Computational uncertainty quantification for the QASPR Silicon circuit prototype	2	09	Jun-09	V&V	SNL
TBD	Evaluation of the impact chip multiprocessors have on SNL application performance	2	09	Sep-09	CSSE	SNL
TBD	Red Storm 284 teraFLOPS upgrade	2	09	Dec-08	CSSE	SNL
TBD	Scalable Analysis Tools for Sensitivity Analysis and UQ	2	09	Jun-09	CSSE	SNL
TBD	Deploy Tripod capabilities for capacity computing environment	2	09	Sep-09	CSSE, FOUS	LLNL, LANL, SNL
TBD	Zia Capability Platform Contract Award	2	09	Jun-09	CSSE	LANL, SNL

Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY10

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	Demonstrate capability to simulate hot and cold environments for safety or primary calculations.	2	10	FY10	IC	LLNL
TBD	Deliver improved AGEX capabilities to users	2	10	Sep-10	IC	LLNL
TBD	Deliver initial capability for shared Monte Carlo transport physics among a suite of ASC codes	2	10	Sep-10	IC	LLNL
TBD	Next generation multiphase, multiscale strength model	2	10	Jun-10	PEM	LLNL
TBD	Numerical convergence study for burn problems	2	10	Sep-10	V&V	LLNL
TBD	Continued Development of Computational Pipeline	2	10		V&V	LLNL
TBD	Release of a Crestone Project code to support high fidelity simulations of a lighting system	2	10	Sep-10	IC	LANL
TBD	Advanced simulation of a lighting system using a Crestone project code	2	10	Sep-10	IC	LANL
TBD	Multiphase EOS assessed in our codes	2	10	Sep-10	PEM	LANL
TBD	Improved mix models in our codes assessed in universal model and in validation suite	2	10	Sep-10	PEM	LANL
TBD	Upgraded PTW strength model for high strain rates	2	10	Sep-10	PEM	LANL
TBD	Implement new non-LTE inline opacity capability in codes	2	10	Sep-10	PEM	LANL
TBD	Advanced HE model for insensitive high explosive	2	10	Sep-10	PEM	LANL
TBD	Assessment of ASC code capabilities for Thermonuclear Applications: Stockpile features and primary outputs	2	10	Jun-10	V&V	LANL
TBD	Assessment of primary burn in an ASC code	2	10	Sep-10	V&V	LANL
TBD	Verification and validation assessment of coupled implicit/explicit codes for simulation of mechanical response in hostile environments	2	10	Sep-10	V&V	LANL
TBD	Comparison of radiation flow treatments using weapons physics calculations on the Roadrunner platform	2	10	Sep-10	V&V	LANL
TBD	Application enablement on next generation platforms	2	10	Sep -10	CSSE	LANL
TBD	Infrastructure Equipment Upgrades Project	2	10	Mar -10	FOUS	LANL
TBD	Roadrunner Phase 3 Transition to Operational Status	2	10	Sep-10	FOUS	LANL
TBD	Coupled turbulent flow, thermal, and quasistatic failure capabilities for assured safety applications	2	10	Sep-10	IC	SNL
TBD	Demonstration of neutron effects on GaAs devices and silicon heterogeneous transistors	2	10	Sep-10	IC	SNL
TBD	Demonstration of advanced algorithms for predicting coupled flow and temperature of fuels in complex environments	2	10	Sep-10	IC	SNL

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	Utilize experimentally validated constitutive model for lead-free solder to simulate aging and reliability of solder joints in stockpile components.	2	10	Sep-10	PEM	SNL
TBD	Computational uncertainty quantification for the QASPR GaAs HBT (Heterojunction Bipolar Transistor) Device Prototype	2	10	Sep-10	V&V	SNL
TBD	Uncertainty analysis for the prediction of the effect of x-ray dose-rate on a circuit or subsystem	2	10	Sep-10	V&V	SNL
TBD	Uncertainty analysis for prediction of cavity SGEMP effects in a reentry body	2	10	Sep-10	V&V	SNL
TBD	Evaluate advanced memory subsystems	2	10	Mar-10	CSSE	SNL
TBD	Zia platform delivery	2	10	Dec-09	CSSE	LANL, SNL
TBD	Zia platform general availability	2	10	Jun-10	CSSE	LANL, SNL
TBD	Zia Platform Application Campaign	2	10	Sep-10	CSSE	LANL, SNL

Detailed Milestone Descriptions for FY09

Milestone (ID#): TLCC clusters in production at LLNL		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Dec-08		
ASC nWBS Subprogram: CSSE, FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone is the culmination of the TLCC platform procurement and will be satisfied when all classified SU are functioning properly on the classified network and science codes are running at scale on the classified side. This is the progression past the acceptance phase of the hardware and software stack on the unclassified side and encompasses running NIF and SSP science at scale on the classified network. Both the 8-SU Juno system for weapons applications and the 2-SU Eos system for NIF applications are part of this milestone.		
Completion Criteria: This milestone is complete when both Juno and Eos are connected to the classified network and are running NIF and SSP science codes with full scheduling and I/O capability.		
Customer: NNSA/ASC HQ		
Milestone Certification Method: Milestone certification will be in the form of written documentation certifying completion of NIF and SSP science runs on the classified network. This documentation will also contain a timeline from system delivery to milestone completion. A hand-off to users is also an implicit part of Milestone completion and will be recognized in that the Eos and Juno systems will be announce as having reached a status of Limited Availability (LA) on the classified network. This announcement to users will be part of the Milestone Certification document.		
Supporting Resources: Fully functioning TOSS as well as fully functioning TLCC hardware are necessary components to meet this milestone. Tri-lab personnel in the form of system administrators and software developers will also be required.		
Supporting Milestones:		
Program	**Title**	**Due Date**
None		
Codes/Simulation Tools Employed: TBD		
Contribution to the ASC Program: Greatly increases productivity of SSP code developers and designers via an order of magnitude increase in number of cycles available for capacity computing.		
Contribution to Stockpile Stewardship: Supports SSP goals including boost studies, UQ analyses, weapons science and predictive capabilities.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	This milestone has dependencies on several external vendors of the TLCC hardware including AMD (processor), Appro (integrator), SuperMicro (motherboard), Wintech (memory) and others not named here.	Medium	Low	Medium

Milestone (ID#): Deploy Dawn ID machine for initial science runs		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Sep-09		
ASC nWBS Subprogram: CSSE, FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone will be a result of work started three years ago with the planning for a multi-petaFLOPS UQ-focused platform (Sequoia) and will be satisfied when a smaller ID version of the final system is delivered, installed, integrated, tested, accepted, and deployed at LLNL for initial science runs in support of SSP mission. The deliverable for this milestone will be a LA petascale computing system (named Dawn) usable for code development and scaling necessary to ensure effective use of a final Sequoia platform (expected in 2011-2012), and for urgent SSP program needs. Allocation and scheduling of Dawn as an LA system will likely be performed informally, similar to what has been used for BlueGene/L. However, provision will be made to allow for dedicated access times for application scaling studies across the entire Dawn resource.		
Completion Criteria: Dawn platform is (1) delivered, (2) installed, integrated and tested, (4) accepted, and (5) deployed at LLNL for initial science runs.		
Customer: SSP early science applications and customers.		
Milestone Certification Method: (1) Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. (2) The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: Computational Systems and Software Environment projects, Facility Operations and User Support projects.		
Supporting Milestones:		
Program	**Title**	**Due Date**
None		
Codes/Simulation Tools Employed: Unclassified IC, PEM or V&V science codes.		
Contribution to the ASC Program: Overall Sequoia platform strategy is based on availability, for at least two or three years, of a significant ID system upon which simulations can evolve and adapt to the architectural evolution inherent in Sequoia platform planning process. This approach enhances the probability of effective use of Sequoia during its lifecycle.		
Contribution to Stockpile Stewardship: In addition to the above strategic approach for simulation code evolvement and adaptation to Sequoia architectures, the Dawn ID system will also be used for urgent nearer-term SSP program needs, in particular for UQ, National Boost Initiative and energy balance PCF pegpost deliverables.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Platform vendor parts commit and build is late, affecting the scheduled delivery timeline	High	Low	Moderate
2	Installation, testing, and/or acceptance is delayed due to unforeseen problems	High	Low	Moderate
3	Initial science applications are not ready to run on Sequoia ID	Moderate	Very Low	Low

Milestone (ID#): Roadrunner system integration readiness		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Jun-09		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: FOUS, IC, PEM		
Description: The focus of this effort will be on completing system hardware deliveries, system installation at Los Alamos, system acceptance testing as contractually required, demonstrated system software, and initial scalability testing of the Roadrunner Phase 3 hybrid system. The Roadrunner statement of work defines the specific requirements to complete the acceptance testing. Included in this effort will be an industry-standard evaluation and LANL specific software acceptance tests. Also included are scientific simulations at scale and the overall system availability over a specific time period. Once the scope of work identified above is completed Roadrunner will be ready to begin on-site integration into the local computing infrastructure, including the network, file systems, archival systems, and system software stack. The Roadrunner system is scheduled to deliver a significantly advanced architecture system that should provide compute power of over a petaFLOPS of computing cycles to the weapons program. The advanced architecture hardware will consist of 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		
Completion Criteria: Requirements in statement of work are met or exceeded, including the hardware and system software. Demonstration of system scalability has also been completed.		
Customer: SSP, to include weapons science customers.		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: Platform funding and system analysts		
Supporting Milestones:		
Program	**Title**	**Due Date**
None		
Codes/Simulation Tools Employed: As specified in the statement of work.		
Contribution to the ASC Program: Brings compute power for ASC program into the petaFLOPS range, enabling scientific calculations in support of the ASC Roadmap.		
Contribution to Stockpile Stewardship: Provides crucial compute cycles and speed to weapon's simulation and physics codes for taking the predictive capability to the next level.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Complete acceptance testing	High	Low	Medium

<table>
<tr><td colspan="4">Milestone (ID#): Performance assessment of major tri-lab ASC Platforms</td></tr>
<tr><td>Level: 2</td><td>Fiscal Year: FY09</td><td colspan="2">DOE Area/Campaign: ASC</td></tr>
<tr><td colspan="4">Completion Date: Sep-09</td></tr>
<tr><td colspan="4">ASC nWBS Subprogram: CSSE</td></tr>
<tr><td colspan="4">Participating Sites: LANL</td></tr>
<tr><td colspan="4">Participating Programs/Campaigns: ASC</td></tr>
<tr><td colspan="4">Description: Assess the performance of all major system installations in the tri-Lab ASC program. Systems to be considered are: full Roadrunner, Dawn, and TLCC. Performance will be done using a realistic ASC application workload. The work will include analyzing the differences between the achieved/measured performance and the achievable/modeled performance on full applications and optimize the system to maximize performance.</td></tr>
<tr><td colspan="4">Completion Criteria: Performance Assessment report completed</td></tr>
<tr><td colspan="4">Customer: WBS 1.5.4.1-LANL-001 System Requirements and Planning</td></tr>
<tr><td colspan="4">Milestone Certification Method:
A program review is conducted and its results are documented.
Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.</td></tr>
<tr><td colspan="4">Supporting Resources: Adequate access to Roadrunner, Dawn, and TLCC. Four FTEs</td></tr>
<tr><td colspan="4">Supporting Milestones:
<table>
<tr><th>Program</th><th>Title</th><th>Due Date</th></tr>
<tr><td>CSSE</td><td>Acceptance of Roadrunner Phase 3 system</td><td>Mar-09</td></tr>
</table></td></tr>
<tr><td colspan="4">Codes/Simulation Tools Employed: Performance Architecture Lab Test Suite. Realistic ASC applications.</td></tr>
<tr><td colspan="4">Contribution to the ASC Program: Enhanced understanding of Roadrunner, Dawn, and TLCC. Ensure that the performance delivered is the achievable one, and pinpoint sources of potential performance degradation.</td></tr>
<tr><td colspan="4">Contribution to Stockpile Stewardship: Enhanced understanding new ASC platforms. Specific source-to-source tools transforming Fortran array syntax and pure and elemental functions</td></tr>
<tr><td colspan="4">
<table>
<tr><th rowspan="2">No.</th><th rowspan="2">Risk Description</th><th colspan="3">Risk Assessment (low, medium, high)</th></tr>
<tr><th>Consequence</th><th>Likelihood</th><th>Exposure</th></tr>
<tr><td>1.</td><td>Loss of personnel</td><td>Medium</td><td>Medium</td><td>Low</td></tr>
</table></td></tr>
</table>

Milestone (ID#): HPSS 7.1 deployment		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Sep-09		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: HPSS release 7.1 will be deployed on the classified and unclassified networks by 4Q09. This release has major metadata performance improvements for small file inserts, tape aggregations for small files, and administrative improvements including real time monitoring of the system.		
Completion Criteria: This milestone is complete when HPSS release 7.1 is installed on both the classified and unclassified network segments and in use by ASC customers.		
Customer: SSP, including weapons science customers.		
Milestone Certification Method: A program review will be conducted and its results are documented. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: HPSS software developers, Operations staff, and stockpile customers		
Supporting Milestones:		

Program	**Title**	**Due Date**
None		

Codes/Simulation Tools Employed: PSI
Contribution to the ASC Program: Faster archive storage processing helps to keep file system user and project areas from running out of space.
Contribution to Stockpile Stewardship: Faster archive storage processing helps to keep file system user and project areas from running out of space.

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1.	HPSS software development schedule slips.	Moderate	Low	Low

<table>
<tr><td colspan="4">Milestone (ID#): Evaluation of the impact chip multiprocessors have on SNL application performance</td></tr>
<tr><td>Level: 2</td><td>Fiscal Year: FY09</td><td colspan="2">DOE Area/Campaign: ASC</td></tr>
<tr><td colspan="4">Completion Date: Sep-09</td></tr>
<tr><td colspan="4">ASC nWBS Subprogram: CSSE</td></tr>
<tr><td colspan="4">Participating Sites: SNL</td></tr>
<tr><td colspan="4">Participating Programs/Campaigns: ASC</td></tr>
<tr><td colspan="4">SNL will investigate the impact of Chip Multi-Processors (CMPs) on the performance of important SNL application codes and the impact of CMPs on the performance and applicability of SNL's system software. This investigation will make algorithmic and architectural recommendations for next generation platform acquisitions, which will be documented in a report.</td></tr>
<tr><td colspan="4">Completion Criteria: Completion of program review and final document published as a SAND report.</td></tr>
<tr><td colspan="4">Customer: NNSA/ASC HQ, CSSE program managers and platform design team members, IC program managers and application developers.</td></tr>
<tr><td colspan="4">Milestone Certification Method:
A program review is conducted and its results are documented.
Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.</td></tr>
<tr><td colspan="4">Supporting Resources: ASC capacity and capability platforms, test beds and early deliverables acquired commercially and via CSSE/Advanced Systems projects.</td></tr>
<tr><td colspan="4">Supporting Milestones: None</td></tr>
<tr><td>Program
N/A</td><td colspan="2">Title</td><td>Due Date</td></tr>
<tr><td colspan="4">Codes/Simulation Tools Employed: This effort leverages the performance analysis and modeling tools developed and/or supported by ASC/CSSE, including but not limited to O|SS and SST.</td></tr>
<tr><td colspan="4">Contribution to the ASC Program: Platform acquisitions will be better equipped to deliver architectures and systems to meet the mission needs of NNSA's computing campaigns.</td></tr>
<tr><td colspan="4">Contribution to Stockpile Stewardship: Platforms</td></tr>
<tr><td rowspan="2">No.</td><td rowspan="2">Risk Description</td><td colspan="3">Risk Assessment (low, medium, high)</td></tr>
<tr><td>Consequence</td><td>Likelihood</td><td>Exposure</td></tr>
<tr><td></td><td>None</td><td></td><td></td><td></td></tr>
</table>

Milestone (ID#): Red Storm 284 teraFLOPS upgrade		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Dec-08		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: SNL will complete the upgrade of the Red Storm computer system to a peak of ~284TF in FY09. The middle section, or 6240 of the 12,960 compute nodes, will be upgraded to 2.2 GHz quad-core Opteron processors. The memory on the entire system will be upgraded to 2 GB per core. The upgrade involves replacing 1560 compute node boards, 6240 compute node processors and all associated memory. The memory from the compute nodes boards that are being replaced will be reused to populate all remaining dual core nodes with 4GB of memory. The existing cabinets, backplanes, interconnect, cabling, and service and I/O nodes will be reused. The system will then have a 38,400 compute node cores with 75TB of compute node memory.		
Completion Criteria: V&V of installed hardware using base level hardware/firmware diagnostics followed by a large-scale application run and analysis of throughput gain over the previous Red Storm configuration.		
Customer: DSW, ASC V&V, ASC Integrated Codes		
Milestone Certification Method: Final report and internal program review presentations of results from the large-scale application run.		
Supporting Resources: SNL FOUS products and personnel		
Supporting Milestones: None		

Program	**Title**	**Due Date**
N/A		

Codes/Simulation Tools Employed: One or more NNSA applications
Contribution to the ASC Program: Progress in meeting dollar and Teraflop Program Goals
Contribution to Stockpile Stewardship: Enhanced Capability for Modeling and Simulation

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	Removal/replacement of modules reduces system stability	Moderate	Low	Medium
2	Additional processor count creates anomalies in system software, reducing productivity/utilization	High	Low	Medium

Milestone (ID#): Scalable Analysis Tools for Sensitivity Analysis and UQ		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Jun-09		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Sensitivity Analysis and UQ require analysis and investigation of ensembles of runs, in which large numbers of runs (anywhere from tens to hundreds) must be analyzed as a group. Tools do not exist to enable investigation of these ensembles of runs, and the sensitivity of the results to changes in the inputs. In practice, analysts cannot examine a significant fraction of the data, and are reduced to examining simple quantities (such as a maximum or minimum of a significant variable) to understand the behavior of the ensemble. Analysis tools developed in this milestone will leverage the scalable scientific and information visualization technologies within ParaView to enable analysts a broader understanding of the ensemble, by providing abstract views of the ensemble, connected to readily-available drill-down to specific data from particular simulation runs.		
Completion Criteria: Delivery preliminary capability to customer that allows interactive exploration of a set of related runs. Demonstrate the use of the capability against representative datasets from customer; however, will provide capability that is widely applicable across codes and problem areas.		
Customer: V&V, Integrated Codes		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a customer is documented.		
Supporting Resources: Datasets delivered by V&V customers		
Supporting Milestones: None		
Program	**Title**	**Due Date**
N/A		
Codes/Simulation Tools Employed: ParaView		
Contribution to the ASC Program: Provide important tools needed to enable V&V/QMU for ASC simulations in support of predictive capabilities. In particular, new capability for sensitivity analysis and UQ.		
Contribution to Stockpile Stewardship: Provide important tools needed to enable V&V/QMU for ASC simulations in support of predictive capabilities. In particular, new capability for sensitivity analysis and UQ.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Inadequate access to V&V and customer 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

Milestone (ID#): Deploy Tripod capabilities for capacity computing environment		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Sep-09		
ASC nWBS Subprogram: CSSE, FOUS		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: Deploy additional Tripod capabilities for capacity computing environment, working towards a responsive and more efficient infrastructure to support computing for QMU and predictivity.		
Completion Criteria: Deploy Tripod capabilities developed in FY08 including: O|SS, Shared Project Work Environment, and Gazebo Test and Analysis Suite. Demonstrate sound software engineering practices with respect to the Tripod Software Stack on the ASC TLCC systems to include, but not be limited to, management of the code repository with tri-lab coordination of the release processes, issue tracking and support. Develop and integrate performance monitoring tools to ensure consistent analysis and reporting of system usage and workload characterization across the tri-labs. The tri-labs will continue to do gap and risk analysis of the Tripod Software Stack; and add new projects, as needed, to address high-priority gaps.		
Customer: ASC HQ; tri-lab program managers		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: CSSE and FOUS		
Supporting Milestones:		

Program	**Title**	**Due Date**
ASC/CSSE & FOUS	TLCC clusters in production at LLNL	June 09

Codes/Simulation Tools Employed: N/A
Contribution to the ASC Program: Common Capacity Computing Environment
Contribution to Stockpile Stewardship: Easy-to-use capacity computing cycles

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	This milestone has dependencies on the on-schedule delivery of the TLCC hardware by system vendors	Medium	Low	Medium

Milestone (ID#): Zia Capability Platform Contract Award		
Level: 2	**Fiscal Year**: FY09	**DOE Area/Campaign**: ASC
Completion Date: Jun-09		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: CSSE		
Description: As a part of the ACES partnership with LANL, develop, issue and evaluate the RFP for the Zia platform and award a contract.		
Completion Criteria: Contract issued to selected vendor		
Customer: NNSA/ASC HQ, tri-lab program managers		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: N/A		
Supporting Milestones: None		

Program	**Title**	**Due Date**
N/A		

Codes/Simulation Tools Employed: N/A
Contribution to the ASC Program: Capability platform deployment
Contribution to Stockpile Stewardship: Capability computing resource to meet demands identified by NNSA/ASC Roadmap

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	NNSA/ASC Platform strategy and budget planning does not fit delivery of a capability platform in FY09.	Low	Low	Low
2	Technology gaps and industry roadmaps do not meet RFP criteria.	Low	Low	Low

Milestone Descriptions for Preliminary FY10

Milestone (ID#): Infrastructure Equipment Upgrades Project		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Mar-10		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: In preparation for the next phase of Supercomputing in 2010 it is necessary to upgrade the existing mechanical and electrical infrastructure in the SCC Facility. The upgrades consists of the procurement and installation of major mechanical equipment (cooling towers, chillers, water cooling skids and air handling units) and major electrical equipment (switchboards and 3000 amp breakers). This milestone will provide the necessary power projected for the Zia Machine in 2010.		

Milestone (ID#): Roadrunner Phase 3 transition to operational status		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Sep-10		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: FOUS		
Description: This effort will culminate in the formal transition of the machine and associated infrastructure to production computing. It will also mark the completion of CD-4 in the Roadrunner CD process. The Roadrunner final system is scheduled to deliver a significantly advanced architecture system that should provide compute power of over a petaFLOPS of computing cycles to the weapons program. The advanced architecture hardware will consist of 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.		

Milestone (ID#): Application enablement on next generation platforms		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Sep-10		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: This project addresses the increasingly complex interface between HW and SW in heterogeneous HPC systems. Given that movement and access of data are the key ingredients to achieving high performance, the emphasis will be on development of tools for memory and communications. The work will help focus the architectural choices and the software environment for the next generation capability machines at the Lab, while having direct applicability to machines currently in use (Roadrunner). This work will focus on a radiation-hydrodynamics workload of interest to ASC. The milestone will provide the ASC community with a suite of memory and communication tools that support the program's application development efforts, assist in developing prototype application software, and enable more accurate system architecture designs.		

Milestone (ID#): Evaluate advanced memory subsystems		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Mar-10		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Develop, with industry and academia partnerships, an advanced memory subsystem to increase the effective performance of SNL application workloads on next generation microprocessors. Advanced functionality includes (but not limited to) the following operations: atomic memory operations, scatter/gather, in-memory copy/zero/fill/etc, and in-memory synchronization. The architecture will be defined, validated and documented with an industrial partner agreement to develop prototypes.		

Milestone (ID#): Zia platform delivery		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Dec-09		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL, LANL		
Participating Programs/Campaigns: ASC		
Description: As a part of the ACES partnership deliver the Zia platform.		

Milestone (ID#): Zia platform general availability		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Jun-10		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL, LANL		
Participating Programs/Campaigns: ASC		
Description: As a part of the ACES partnership deploy the Zia platform.		

Milestone (ID#): Zia Platform Application Campaign		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: Sep-10		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL, LANL		
Participating Programs/Campaigns: ASC		
Description: An initial capability campaign using a modern ASC code that targets the capabilities of the Zia platform design.		

VI. ASC Roadmap Drivers for FY09–FY10

Table VI-1. ASC Roadmap Drivers for FY09-10[6]

Focus Area 1	
2009	Modular Physics and Engineering Packages For National Weapons Codes
Focus Area 2	
2009	Science-Based Replacement for Knob (Ad Hoc Model) #1
2010	Science-Based Models for Neutron Tube Simulations
Focus Area 3	
2009	Shared Weapons Physical Databases
2012	Uncertainty Quantification (UQ) Methodology for QMU
Focus Area 4	
2009	Petascale Computing

[6] The ASC Top Ten Risks table was originally published in the *ASC Program Plan FY05*.

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[7]

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.

[7] 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 (LEP, 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 and engineering model development
6	Data not available for input to new physics models or for model validation.	High	Moderate	MEDIUM	Work with Science and Engineering 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 Campaign Annual Performance Results and Targets

(R = Results; T = Target)

Performance Indicators	FY 2005 Results	FY 2006 Results	FY 2007 Results	FY 2008	FY 2009	FY 2010	FY 2011	FY 2012	FY 2013	Endpoint Target
Strategic Goal 2.1 (Nuclear Deterrent) GPRA Unit Program Goal 2.1.30.00, Advanced Simulation and Computing Campaign										
Adoption of ASC Modern Codes: The cumulative percentage of simulation runs that utilize modern ASC-developed codes on ASC computing platforms as measured against the total of legacy and ASC codes used for stockpile stewardship activities (Long-term Outcome)	N/A	R: 50%	R: 63%	T: 72%	T: 80%	T: 85%	T: 90%	T: 95%	T: 100%	By 2013, ASC-developed modern codes are used for all simulations on ASC platforms. Adoption of Modern ASC Codes will enable a responsive simulation capability for the nuclear weapons complex. This measure is meant to show how quickly ASC codes are being adopted by the user community in place of legacy codes.
Reduced Reliance on Calibration: The cumulative percentage reduction in the use of calibration "knobs" to successfully simulate nuclear weapons performance (Long-term Outcome)	N/A	R: 2%	R: 8%	T: 16%	T: 25%	T: 33%	T: 41%	T: 50%	T: 58%	By 2018, the four major calibration knobs affecting weapons performance simulation have been replaced by science-based, predictive phenomenological models. Reduced reliance on calibration will ensure the development of robust ASC simulation tools, These tools are intended to enable the understanding of the complex behaviors and effect of nuclear weapons, now and into the future, without nuclear testing.
ASC Impact on SFI Closure: The cumulative percentage of nuclear weapon Significant Finding Investigations (SFIs) resolved through the use of modern (non-legacy) ASC codes, measured against all codes used for SFI	N/A	R: 10%	R: 25%	T: 37%	T: 50%	T: 62%	T: 75%	T: 87%	T: 100%	By 2013, ASC codes will be the principal tools for resolution of all SFIs. This demonstrates how valuable the ASC tools are for meeting the needs of the weapon designer's analysts by documenting the impact on closing SFIs.

Performance Indicators	FY 2005 Results	FY 2006 Results	FY 2007 Results	FY 2008	FY 2009	FY 2010	FY 2011	FY 2012	FY 2013	Endpoint Target
resolution (Long-term Outcome)										
Code Efficiency: The cumulative percentage of simulation turnaround time reduced while using modern ASC codes (Efficiency)	N/A	R: 6%	R: 7%	T: 13%	T: 26%	T: 32%	T: 39%	T: 45%	T: 50%	By 2013, achieve a 50% reduction in turnaround time, as measured by a series of benchmark calculations, for the most heavily used ASC codes. To show code efficiency by demonstrating that simulation time decreases as the ASC codes mature.
NOTE: Performance measures were revised in 2007 to be consistent with new program roadmap.										

Appendix A. Glossary

3D	Three Dimensional
ACES	Alliance for Computing at Extreme Scale
ASAP	Academic Strategic Alliance Program
ASC	Advanced Simulation and Computing
ASC	Accelerated Strategic Computing Initiative
AST	Alliance Strategy Team
C-SAFE	Center for Simulation of Accidental Fires and Explosions
CAD	Computer Aided Design
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CCF	Central Computing Facility
CEI	Computational Engineering International
CRASH	Center for Radiative Shock Hydrodynamics
CRT	Computer Resource Team
CSAR	Center for Simulation of Advanced Rockets
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
DoD	Department of Defense
DOE	Department of Energy
DSW	Directed Stockpile Work
DTRA	Defense Threat Reduction Agency
ESN	Enterprise Secure Net
FOUS	Facility Operations and User Support (WBS 1.5.5)
FVM	Finite Volume Method
GA	General Availability
GB	Gigabytes
GB/sec.	Gigabytes per Second
GMV	General Mesh Viewer
GPFS	Global Parallel File System
HTGL	High Temperature Gasdynamics Laboratory
GPU	Graphical Processing Units
HPC	High-Performance Computing

HERT	SNL's HPC Estimations and Requirements Tool
HPSS	High-Performance Storage System
HQ	ASC Headquarters
I/O	Input/Output
IAA	Institute for Advanced Architectures and Algorithms
ID	Initial Delivery
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LDCC	Laboratory Data Communications Center
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
MPI	Message Passing Interface
MPM	Material Point Method
MSC	Multidiscipline Simulation Center
MURI	Multidisciplinary University Research
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
O|SS	Open|SpeedShop
OCF	Open Computing Facility
OLN	Open LabNet
ONR	Office of Naval Research
PECOS	Center for Predictive Engineering and Computational Sciences
PRISM	Center for Prediction of Reliability, Integrity, and Survivability of Microsystems
PSAAP	Predictive Science Academic Alliance Program
PSI	Parallel Storage Interface
QMU	Quantification of Margins and Uncertainties
R&D	Research and Development
RH	Radiation Hydrodynamics
RAS	Reliability, Availability, and Serviceability
RDMA	Remote Direct Memory Access

RSRM	Reusable Solid Rocket Motor
SCC	Nicolas C. Metropolis Center for Modeling and Simulation
SCF	Secure Computing Facility
SDM	Scientific Data Management
SFI	Significant Finding Investigation
SLA	Service Level Agreement
SNL	Sandia National Laboratories
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STS	Stockpile-to-Target Sequence
SU	Scalable Unit(s)
SWL	Shock Wave Lens
TLCC	Tri-Lab Linux Capacity Cluster
TLCC07	Tri-Lab Linux Capacity Cluster for 2007
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
TST	Tri-Lab Sponsor Team
UQ	Uncertainty Quantification
V&V	Verification and Validation
VTF	Virtual Test Facility
WAN	Wide Area Network
ZFS	Zettabyte File System

Appendix C. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.1	Capability Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manual Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.2	Capacity Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manual Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.3	Advanced Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.4	System Software and Tools	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Jim Ang, SNL, 505-844-0068, jaang@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Jim Ang, SNL, 505-844-0068, jaang@sandia.gov
1.5.4.6	Post-Processing Environments	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.4.7	Common Computing Environment	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Judy Sturtevant, SNL, 505-845-9448, jesturt@sandia.gov
1.5.5	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Zepper, SNL, 505-845-8421, jdzeppe@sandia.gov
1.5.5.1	Facilities, Operations, and Communications	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov

WBS	Title	Contact
1.5.5.2	User Support Services	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.5.3	Collaborations	Lynn Kissel, LLNL, 925-423-7940, kissel1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Martin Pilch, SNL, 505-845-3047, mpilch@sandia.gov Paul Yarrington, SNL, 505-844-7504, pyarrin@sandia.gov
1.5.1.4-TRI-001	Caltech, Center for Simulating Dynamic Response of Materials	Michael Ortiz, 626-395-4530, Ortiz@aeor.caltech.edu Dan Meiron, 626-395-3424, dim@its.caltech.edu Michael Aivazis, 626-395-3424, aivazis@caltech.edu Paul Dimotakis, 626-395-6811 Bill Goddard, 626-395-2731, wag@wag.caltech.edu Michael Ortiz, 626-395-4530, ortiz@aeor.caltech.edu Dale Pullin, 626-395-6081, dale@galcit.caltech.edu
1.5.1.4-TRI-001	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, moin@stanford.edu Juan Alonso, 650-723-9954, jjalonso@stanford.edu Heinz Pitsch, 650-736-1995, h.pitsch@stanford.edu William Dally, 650-725-8945, billd@csl.stanford.edu Pat Hanrahan, 650-723-8530, hanrahan@cs.stanford.edu
1.5.1.4-TRI-001	University of Chicago, Center for Astrophysical Flash Phenomena	Donald Q. Lamb, 773-702-7194, d-lamb@uchicago.edu Todd Dupont, 773-702-3485, Dupont@cs.uchicago.edu Rusty Lusk, 630-252-7852, lusk@mcs.anl.gov Tomasz Plewa, 773-834-3227, tomek@flash.uchicago.edu Andrew Siegel, 773-834-8501, siegela@flash.uchicago.edu Jim Truran, 773-702-9584, truran@nova.uchicago.edu
1.5.1.4-TRI-001	University of Illinois, Center for Simulation of Advanced Rockets	Michael T. Heath, 217-333-6268, m-heath@uiuc.edu William A. Dick, 217, 244-7235, wdick@uiuc.edu Robert A. Fiedler, 217-333-3247, rfiedler@uiuc.edu
1.5.1.4-TRI-001	University of Utah, Center for Simulation of Accidental Fires and Explosions	David W. Pershing, 801-581-5057, David.Pershing@utah.edu Chuck Wight, 801-581-8796, Chuck.Wight@utah.edu Tom Henderson, 801-581-3601, tch@cs.utah.edu Philip Smith, 801-585-3129, smith@crsim.utah.edu Patrick McMurtry, 801-581-3889, mcmurtry@eng.utah.edu Eric Eddings, 801-585-3931, eddings@che.utah.edu
1.5.1.4-TRI-001	Purdue University, Center for Prediction of Reliability, Integrity and Survivability of Microsystems	Jayathi Murthy, 765-494-5701, jmurthy@ecn.purdue.edu

WBS	Title	Contact
1.5.1.4-TRI-001	University of Michigan, Center for Radiative Shock Hydrodynamics	Paul Drake, 734-763-4072, rpdrake@umich.edu
1.5.1.4-TRI-001	University of Texas, at Austin, Center for Predictive Engineering and Computational Sciences	Bob Moser, 512-471-0093, rmoser@mail.utexas.edu

Appendix D.
WBS 1.5.1.4-TRI-001 Academic Alliance Centers

The Academic Alliance Centers project includes research activities at the eight funded academic centers as part of the ASAP and PSAAP, as listed below.

Academic Strategic Alliance Program:

- University of Chicago
- University of Illinois at Urbana-Champaign (UIUC)
- University of Utah

Predictive Science Academic Alliance Program:

- California Institute of Technology (Caltech)
- Purdue University
- Stanford University
- University of Michigan
- University of Texas, at Austin

The ASAP Centers will focus on bringing to conclusion the work they had pursued under the earlier phase of the Alliances program to demonstrate the utility of large scale computational simulations for solving challenging multi-physics applications on ASC HPC platforms.

The new PSAAP Centers will focus on the emerging field of predictive science, i.e. the application of verified and validated computational simulations to predict the behavior of complex systems where routine experiments are not feasible. The centers will focus on unclassified applications of interest to NNSA and its three national laboratories: Lawrence Livermore National Laboratory, Los Alamos National Laboratory and Sandia National Laboratories. The PSAAP centers will develop not only the science and engineering models and software for their large-scale simulations, but also methods associated with the emerging disciplines of V&V and UQ. The goal of these emerging disciplines is to enable scientists to make precise statements about the degree of confidence they have in their simulation-based predictions.

Both the ASAP and PSAAP Centers will continue to help develop for NNSA and the Nation the workforce of the future, wherein simulations will be pervasive and instrumental in important, high-impact decision-making processes.

In addition, the project also funds the technical management support services for these Centers by the tri-lab Alliance Strategy Team (AST), the Computer Resource Team (CRT) and Tri-Lab Sponsor Teams (TSTs). The AST manages and coordinates the project among the labs, the 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 annual

spring TST meetings with the Centers. The CRT is responsible for assisting the Centers' users in accessing and utilizing the ASC computing resources at the 3 NNSA Labs.

Planned activities in FY09:

- Carry out research activities as described in the section below.
- Maintain active interactions with the ASAP and PSAAP Centers, review all technical work, and provide all required technical, HPC access and management support to the Centers.

Expected deliverables in FY09:

- Efficient performance of FLASH on multi-core platforms (U of Chicago).
- Simulations of RSRM to exercise the then-most recent features of Rocstar (UIUC).
- Code validation using scaled thermal explosion data (U of Utah).
- Full-system Ta/Ta ballistic runs for UQ analysis, including verification (nonlinear sensitivity analysis) and validation (in coordination with full system experiments) runs (Caltech).
- Results from V&V simulations for MPM and FVM, both separately and coupled, as well as V&V for LAMMPS molecular dynamics simulations (Purdue).
- Development of a low-fidelity computational infrastructure for the simulation of the HyShot flight (Stanford).
- A software report detailing the SWMF framework modified to jointly run BATSRUS (hydrodynamics) and PDT (radiation transport) (U of Michigan).
- Completion of first coupled simulations of reentry capsule with chemistry, turbulence, radiation and ablation models including sensitivity analysis of heat fluxes and ablation rates to models and model parameters (U of Texas).

Preliminary planned activities in FY10:

- Continue to conduct planned research activities at the eight Alliance Centers
- Organize a technical conference focusing on Predictive Science with other federal agencies.

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."

The FLASH code is a fully modular, extensible, community code that is capable of simulating a wide variety of problems is 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.

Planned activities and expected deliverables in FY09:

- Enable FLASH to perform efficiently on multi-core platforms
- Develop a Petascale scientific analysis pipeline and archive
- Validate Type Ia supernovae (SNe Ia); the Flash Center will expand its large-scale, 3D simulations of SNe Ia from the "gravitationally confined detonation" model to other current models, including the pure deflagration model and the "deflagration to detonation" model

Expected deliverables in FY09:

- Modeling of FLASH performance on next-generation machines in order to prepare FLASH to perform efficiently and scale well on them
- Parallelization of SEDONA Type Ia supernova radiation transport code
- Completion of verification studies of buoyancy-driven turbulent nuclear burning

Preliminary planned activities in FY10:

- Complete effort to thread FLASH
- Parallelize the Phoenix Type Ia supernova (SN Ia) radiation transport code
- Extend its large-scale, 3D simulations of SNe Ia to new initial conditions, including the effects of rotation

University of Illinois at Urbana-Champaign

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.

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, Rocinteract), 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 (Roccom, Rocman, Rocprop, Rocface, and others).

Planned activities in FY09:

- Develop design acquisition and preliminary simulations for NASA crew launch abort motor
- Initially implement non-dissipative unstructured fluid dynamics solver (Rocflu-ND)

Expected deliverables in FY09:

- Simulations of RSRM to exercise the then-most recent features of Rocstar
- Simulations of NASA booster separation motor
- Chimera mesh technology for turbulence and other high-fidelity simulation needs (for example, acoustics)

Preliminary planned activities in FY10:

- Validate Rocflu-ND for SRM applications
- Study acoustic effects of 3D vortex shedding in RSRM simulations downstream of intersegment inhibitors
- Understand better the effects of propellant heterogeneity on motor performance
- Implement and validate simulation modules for thermal and mechanical insult for propellants and other energetic materials
- Extend Rocstar simulation technology to applications beyond solid propellant rockets

University of Utah

Center for Simulation of Accidental Fires and Explosions

Since 1997, the University of Utah Center for Simulation of Accidental Fires and Explosions (C-SAFE) has been building a state-of-the-art science-based high-performance simulation code to predict the explosive response of a steel container filled with PBX9501 embedded in a large jet fuel fire. The simulations involve a broad range of conditions involving characteristics of the container, its location relative to the fire, and the influence of wind direction and intensity.

Planned activities in FY09:

- Extend the validation and UQ of the heat flux from fires to produce uncertainty ranges (error bars) based on consistency between the simulation results and the experimental data from Sandia National Laboratory, the University of Nevada at Reno, and ATK.
- Carry out an analysis of the simulation data from 16 scenarios to characterize the violence associated with each of the scenarios. Preliminary analysis indicates that the expected qualitative behavior of higher violence with lower heating rates is observed, but in order for these observations to be meaningful, a quantitative

analysis must be done that considers velocity of the container walls, container breakup and rate of release of high pressure product gases.

- Begin simulations of packed arrays of explosives, along with further model development, which we anticipate will enable us to capture the deflagration to detonation transition.

Expected deliverables in FY09:

- Validation using scaled thermal explosion data.
- Development of first-principles based EOS and constitutive model for PBX-9501 that more accurately accounts for the influence of temperature and pressure on the mechanical response of the PBX.
- Validation Simulations for the Coupled Heat-Up/Explosion Scenario.
- Experimental data in a form suitable for model validation will be provided on interferometry in He plumes, LIF and Raman spectroscopy for gas species in a simple hydrocarbon flame, and descriptions of physical and thermal properties of soot deposits on containers in flames as a function of flame conditions.

Preliminary planned activities in FY10:

- Complete the UQ for predicted heat flux to containers in or near jet fuel pool fires to include the error budget to quantify the error sources in the predictions.
- Publish results of numerical analysis findings and results from simulation science.
- Complete an experimental combustion study using interferometry for hydrocarbon flames, and LIF and Raman spectroscopy for a flame with a complex hydrocarbon mixture.
- Provide completed documentation and archiving, in conjunction with an associated NSF project, to ensure the final code can be used by a broader user base.

California Institute of Technology

The Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials

Caltech's Multidiscipline Simulation Center (MSC) overarching objective is the development of a multidisciplinary Predictive Science methodology focusing on high-energy-density dynamic response of materials and the demonstration of the methodology by means of a concerted and highly integrated experimental, computational, and analytical effort focusing on an overarching ASC-class problem: hypervelocity normal and oblique impact of projectiles on metallic and non-metallic targets, at velocities up to 10 km/s. Hypervelocity impact gives rise to pressures in the Mbar range and strain-rates up to 1011s-1, providing a grand-challenge problem in Predictive Science that is also well-matched to the direct interests of the NNSA mission. The overarching hypervelocity impact application, in conjunction with a rigorous and novel methodology for model-based UQ, will provide the intellectual backbone of the Center and its chief organizing principle. In particular, the QMU will drive and closely coordinate the experimental, computational, modeling, software development, V&V efforts within a yearly assessment format.

Planned activities in FY09:

- Verify and perform analysis of our Lagrangian finite element capability for the simulation of oblique ballistic impact of single Ta plate by Ta projectile in the 2-3 km/s impact velocity range.
- Apply the verified and optimized Lagrangian finite element capability, in close coordination with ballistic experiments to be conducted at Caltech's Small Particle Hypervelocity Range (SPHR), to the quantification of uncertainties in the ballistic performance of Ta plates.

Expected deliverables in FY09:

- Determination of UQ analysis, in terms of confidence factors, on the extent to which the ballistic performance of the plates can be predicted by our simulation capability.
- Full-system Ta/Ta ballistic runs for UQ analysis, including verification (nonlinear sensitivity analysis) and validation (in coordination with full system experiments) runs.
- Full-system Ta/Ta ballistic data in support of UQ analysis.
- Development of improved concentration-of-measure inequalities for UQ analysis that self-adapt to the nonlinear structure of the response function.
- Implementation of massively parallel optimization algorithms within the VTF that support UQ analysis and make effective utilization of NNSA petascale computing resources.
- Development and validation of a fast multiscale model of polycrystalline behavior for Ta.
- Extension and validation of Ta EoS to higher pressures and temperatures.
- Integration of new strength EoS and transport models into the VTF.
- Split Hopkinson (Kolsky) pressure bar experiments with shear compression specimen and Shock Wave Lens experiments in support of strength and EoS model development.
- ReaxFF/TB molecular dynamics runs of dense Ta plasmas and determination of initial conditions for continuum plasma simulations.

Preliminary planned activities in FY10

- Run Eulerian hydrocode simulation of oblique hypervelocity impact of single Ta plate by Ta projectile in the 5-10 km/s impact velocity range.
- Implement full-system Ta/Ta hypervelocity runs for UQ analysis, including verification (nonlinear sensitivity analysis) and validation (in coordination with full system experiments) runs.
- Develop improved concentration-of-measure uncertainty bounds that exploit the hierarchical and multiscale structure of hypervelocity.
- Implement an improved concentration-of-measure uncertainty bounds within the VTF that makes effective utilization of NNSA petascale computing resources.
- Develop and validate a fast multiscale model of polycrystalline behavior for Fe including solid-solid phase transitions.
- Study thermal and electrical conductivities and optical frequency response function of the dense plasma state of Ta for a wide range of temperatures and densities.

Purdue

Center for Prediction of Reliability, Integrity and Survivability of Microsystems

The overall objective of the Center for Prediction of Reliability, Integrity and Survivability of Microsystems (PRISM) is to accelerate substantially the integration of MEMS technologies into NNSA stockpile monitoring and weapons systems. PRISM aims to significantly improve our understanding of the long-term reliability of MEMS and their survivability in harsh environments by simulating rigorously, and at multiple scales, the physics of failure, accounting for the coupled electrical, mechanical, thermal and materials behavior of MEMS, from atoms to devices. Advanced simulation software developed by the Center will be encapsulated in an integrated simulation system, MEMOSA (MEMS Overall Simulation Administrator).

Planned activities in FY09:

- Develop MEMOSA to address MEMS structural response under the action of fluid forces. Structural response will be addressed using the material point method (MPM), while the fluid will be addressed using the finite volume method (FVM). The development of the FVM and MPM solvers will be initiated based on existing solvers and methodologies, enhanced to address MEMS applications, and integrated to operate in a coupled fashion.
- Develop a V&V suite for the LAMMPS molecular dynamics code and complete a suite of V&V. Simulations of metal-metal and metal-dielectric contact will be performed using the potentials already existing in LAMMPS.

Expected Deliverables in FY09:

- Results of structural response simulations of MEMS-related structures using the material point method under inertial and impulse actuation.
- Results of fluid flow simulations in MEMS-related geometries under prescribed solid motion.
- Results of fluid-structure interaction simulations in typical MEMS structures.
- Data from metal-metal and metal-dielectric contact simulations using LAMMPS molecular dynamics code with existing interatomic potentials and software.
- Data from surface reaction simulations using reactive force field molecular dynamics.
- Results from V&V simulations for MPM and FVM, both separately and coupled, as well as V&V for LAMMPS molecular dynamics simulations.
- Experimental data for dynamic response of MEMS-related structures under inertial and impulse excitation.
- Microstructure characterization data for MEMS-related structures.

Preliminary planned activities in FY10

- Continue the development of MEMOSA to address (i) structural response under the action of fluidic and electrostatic forces in the presence of thermal gradients, (ii) modeling and simulation of dielectric charging, and (iii) enhancement of molecular dynamics simulations of metal-dielectric contact.
- Simulate integrated fluid-structure-electrostatic-thermal response of MEMS-related structures.

- Perform quantification of uncertainty based on collocation using FVM and MPM solver modules for design of experiments.
- Parallelize MEMOSA solvers.
- Develop first-principles based force-fields for molecular dynamics and simulations using these force fields to study metal-dielectric contact.
- Perform V&V of models related to electrostatics, thermals and dielectric charging.
- Complete dielectric charging and electro-thermal experiments to provide validation data.

Stanford University

The Center for Predictive Simulations of Multi-Physics Flow Phenomena with Application to Integrated Hypersonic Systems

The objective of the Center is to enable predictive multi-physics simulations of a fully integrated hypersonic vehicle. Its overarching problem is the simulation of air-breathing hypersonic vehicles with a special focus on the prediction of off-design, transient conditions and their associated failure modes. Air-breathing hypersonic vehicles are envisioned as a means for reliable low-cost access to space. These vehicles are highly integrated systems whose performance depends on complex physics and the interactions between all of its components. Such performance-critical systems cannot be predicted with today's state-of-the-art simulation capabilities: a radically new integrated approach is required.

World-class experimental facilities in Stanford's High Temperature Gasdynamics Laboratory (HTGL) will be used to conduct validation experiments for the key component physics and models. Verification methods will be developed and implemented as an integral part of our effort at both the component and the system levels. We will leverage advanced computer science methods developed at Stanford to ensure scalability, program correctness, and portability to future platforms with very large numbers of cores.

Planned activities in FY09:

- Implement an efficient UQ technique to enable the evaluation of the confidence in the simulations
- Validate the predictions with respect to the available flight data (received from U. Queensland in March 08).

Expected deliverables in FY09:

- Development of a low-fidelity computational infrastructure for the simulation of the HyShot flight. The physical processes that will be modeled include: strong shocks and non-equilibrium gas-dynamics effects at the vehicle nose; transition to turbulence triggered by trips; Reynolds-averaged turbulence models; inlet/isolator shock train; combustion modeling for hydrogen fuel; and heat transfer control

Preliminary planned activities in FY10

- Initiate in-house experimental campaign: fuel mixing in supersonic stream; shock tube experiments to obtain reaction rates for advanced fuels; and plasma control
- Run simulations using PSAAP computational infrastructure to support experimental design and planning

- Implement LES with low-dissipation shock capturing and explicit LES filtering
- Develop a low-fidelity heat release model for hydrogen combustion and reduced mechanisms for advanced fuels
- Begin integration effort for stable and accurate non-matching coupling of the conjugate heat transfer to aerodynamic codes, and validation and verification
- Develop UQ methodology for ROM and Likelihood-averaging approaches and shock/boundary layer interaction (Monte Carlo experiment)
- Continue domain-specific language development to support a subset of the computational infrastructure for efficient multi-core implementation

University of Michigan

The Center for Radiative Shock Hydrodynamics

The Center for Radiative Shock Hydrodynamics (CRASH) is advancing predictive science in the nationally important area of radiation hydrodynamics (RH) via a unified, multi-prong approach. To substantially improve the ability to do predictive simulations of high-energy-density and astrophysical flows, Center researchers are:

- Developing a software framework for RH to serve as a testbed for development, V&V of RH modeling elements.
- Developing a system for hierarchically validating the software framework.
- Extending an existing experimental effort, centered on radiative shocks, to obtain data and quantify uncertainties in the experiments.
- Simulating these experiments and quantifying the accuracy of the simulations.
- Establishing a doctoral program track for Predictive Science and Engineering.

Planned activities in FY09:

- Deliver report of a simulation of the signature experiment with version 1 of the CRASH code
- Attempt to create identical initial conditions and to repeat an identical set of measurements as many times as is feasible during the experimental day, to assess experimental variability
- Modify the BATSRUS hydrodynamics to: i) solve the hydrodynamic equations described in our proposal, incorporating simple models for equation of state and ionization, ii) to use a first-order level-set algorithm to track material interfaces, and iii) as required to link through SWMF to the radiation transport calculation
- Develop in the PDT radiation transport code: i) a gray transport solver working with tabular opacities or fits, and ii) the software required to accomplish the needed electron-temperature iterations for the thermal source terms and iii) the software to link through SWMF to BATSRUS

Expected deliverables in FY09:

- A software report detailing the SWMF framework modified to jointly run BATSRUS (hydrodynamics) and PDT (radiation transport)
- A UQ report, describing: an operating Bayesian inference engine, adapted from existing methods, and designed to allow the use of either uncertainty intervals or distributions and to retain the data necessary to quantify the sensitivity to initial

assumptions; a functioning suite of first-year code-verification tests; and a functioning solution-verification test suite for hydrodynamics, interface tracking, and radiation transport

Preliminary planned activities in FY10:

- Develop improved capabilities for the simulation, including: ensuring that SWMF, BATSRUS, and PDT are able to run as a common system that can be called by an uncertainty evaluator; advancing BATSRUS to include i) electron heat conduction, ii) second-order level set tracking of material interfaces, iii) tabular EOS, and iv) the use of steady-state adjoint methods for grid adaptation; and extending the radiation transport calculation to include i) multigroup diffusion based on multigroup, tabular opacities, ii) discrete-ordinates radiation transport by means of PDT.
- Carry out the Year 2 Experiment, which is expected to focus on the initial evolution of the experiment, to improve the uncertainty assessment of the initial conditions for the codes run under SWMF.
- Develop first-generation tools for assessment of predictive capability, using the first version of the Bayesian emulator.

University of Texas

The Center for Predictive Engineering and Computational Sciences (PECOS)

The goal of the PECOS Center is to develop the next generation of advanced computational methods for predictive simulation of multi-scale, multi-physics phenomena relevant to the NNSA, and to apply these methods to the problem of reentry of vehicles into the atmosphere.

Simulation of vehicle reentry into the atmosphere requires modeling of the interaction of extremely high temperature gas flows with the high temperature response of materials, in particular the vehicle's thermal protection system. The high gas temperatures produce chemical dissociation, thermal non-equilibrium, and possibly ionization. Radiative heat transfer is an important part of the heat load on the vehicle, while transition and turbulence greatly enhance the rate of heat transfer, as does the transient firing of reaction control jets used to control vehicle attitude. During reentry, the ablative thermal protection system responds viapyrolysis, chemical reaction, and formation and mechanical degradation of a refractory char layer. Models of these high-energy, multiscale, multiphysics phenomena will be integrated into a unified simulation code, designed to support predictive simulation.

Planned activities in FY09:

- Implement model development and code verification: 1) Ablation Modeling (implementation of CMA-type locally-one-dimensional ablation models for integration with HyFlow; implementation and evaluation of existing thermo-mechanical TPS constitutive model; and collection of experimental morphological and thermogravimetric data); 2) Aerothermochemistry (evaluation of turbulent reaction model in DPLR/US3D, and assessment of validity of current multi-temperature rate models); 3) Hypersonic Flow code (HyFlow) (rigorous code verification of DPLR/US3D, and integration of ablation and radiation models using DPLR/US3D); and 4) Radiation
- Develop and implement delayed rejection adaptive Markov chain Monte Carlo sampling algorithms for Bayesian inversion; development and implementation of

Monte Carlo sampling algorithms for uncertainty propagation; and development of Physics-Aware Interface Definition coupling framework

- Implement calibration and validation of turbulence models in DPLR/US3D using available legacy data

Expected deliverables in FY09:

- Implementation of simplified P1 radiation model: turbulence (implementation and verification of the high speed boundary layer DNS code, completion of initial DNS simulations, and collection of experimental validation data for Mach 5 turbulent boundary and interaction with a shock)
- Completion of first coupled simulations of reentry capsule with chemistry, turbulence, radiation and ablation models including sensitivity analysis of heat fluxes and ablation rates to models and model parameters

Preliminary planned activities in FY10:

- Continue working on the model development, implementation and code verification milestones in cooperative contracts
- Develop and implement non-intrusive reduction methods into MCMC sampling for Bayesian inversion
- Develop and implement non-intrusive reduction methods into uncertainty propagation software
- Develop adjoint ablation algorithms and code
- Calibrate and validate a CMA-type ablation model using thermogravimetric experiments
- Conduct coupled system simulations of the reentry capsule using new PAID infrastructure and uncertainty propagation including first quantification of uncertainty in heat fluxes and ablation rates using available estimates of input uncertainties

Appendix E.
ASC Obligation/Cost Plan

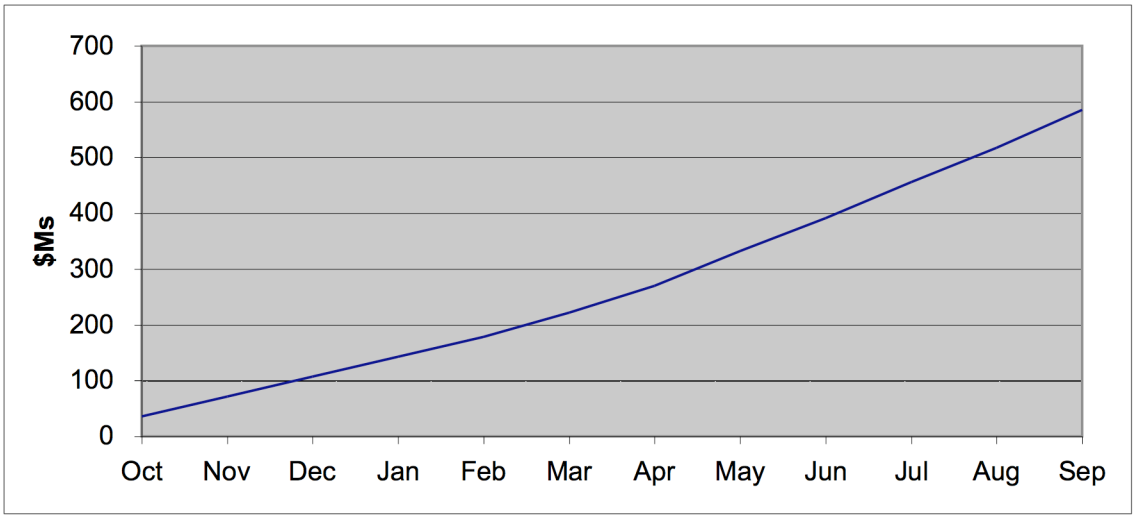

Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
35.523	71.548	107.472	142.233	177.871	222.056	269.830	331.930	391.497	455.713	517.286	585.738

Figure D-1. ASC obligation/cost plan for FY08.

Appendix F.
ASC Performance Measurement Data for FY09

Table E-1. ASC Performance Measurement Data for FY09

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provide the computational science and computer simulation tools necessary for understanding various behaviors and effects of nuclear weapons for responsive application to a diverse stockpile and scenarios of national security.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY07	**FY08**	**FY09**	**FY10**	**FY11**	**FY12**	**FY13**	
CODE PREDICTIVITY: Biennial progress toward a desired end-state through a series of incremental targets of code usability and reliability for ASC applications as measured by the Code Maturity Index.	Baseline	TBD	TBD	TBD	TBD	TBD	TBD	By 2015, achieve 100% of target code maturity as assessed using the ASC Code Predictivity Characteristic Matrix and measured by the Code Maturity Index.
STOCKPILE IMPACT: The cumulative percentage of high consequence stockpile activities to which the modern ASC codes are the predominant simulation tools as measured by the national Stockpile Application Index (nSAI).	Baseline	55%	65%	77%	88%	100%	100%	By 2012, designers and analysts will use ASC codes predominately in 100% of the defined list of high consequence stockpile activities as measured by the nSAI.

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provide the computational science and computer simulation tools necessary for understanding various behaviors and effects of nuclear weapons for responsive application to a diverse stockpile and scenarios of national security.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY07	**FY08**	**FY09**	**FY10**	**FY11**	**FY12**	**FY13**	
CAPABILITY COMPUTING: The annual percentage of usage of ASC capability platforms for simulations that use at least 30% of the platform capability, as defined by total available node hours.	Baseline	18%	25%	32%	39%	45%	45%	By 2012, use 45% of the total each ASC capability platform for simulations that use at least 30% of the platform capability.
COST EFFICIENCY: The annual cost ($Ms) per petaFLOPS (quadrillion floating point operations per second) to procure, upgrade, operate, manage and maintain ASC computing platforms.	Baseline	1,089	762	534	320	192	121	By 2015, achieve a decrease in the life-cycle cost per petaFLOPS to $100M from FY2007 Baseline.
Code Predictivity milestones	Deliver a physics-based sub-grid model to support energy balance resolution (LLNL) Enhanced Pu multiphase Equation of State capability (LANL) Predictive failure capabilities in SIERRA mechanics (SNL) Fundamental simulations of material response and plasma physics on HPC platforms (LANL)							
Stockpile Impact milestones	ASC simulations supporting the National Technical Nuclear Forensics Attribution program (LANL) Assessment of nuclear physics uncertainties (LLNL) Aged temperature/radiation aware model of Complementary Metal Oxide Semiconductor integrated circuit technology (SNL)							
Capability Computing milestones	Infrastructure Plan for ASC Petascale Environments (LLNL, LANL, SNL)							

Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550-808