FY10-FY11 Implementation Plan

Volume 2, Rev. 0.5

Advanced Simulation & Computing™

September 24, 2009

NA-ASC-119R-08-Vol. 2, Rev. 0.5-IP

LLNL-TR-416988

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

FY10–11 IMPLEMENTATION PLAN
Volume 2, Rev. 0.5

September 24, 2009

Approved by:

Robert Meisner, NNSA ASC Program Director (acting)	Signature	*9/24/09* Date
James Peery, SNL ASC Executive	Signature	*9/10/09* Date
Michel McCoy, LLNL ASC Executive	Signature	*9/22/09* Date
John Hopson, LANL ASC Executive	Signature	*9/17/09* Date

ASC Focal Point
Robert Meisner
NA 121.2
Tele.: 202-586-0908
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
I.	Executive Summary		
II.	Introduction		
III.	Accomplishments		
IV.	Product Descriptions		
V.	ASC Level 1 and 2 Milestones		
VI.	ASC Roadmap Drivers for FY10–FY11		
VII.	ASC Risk Management		
VIII.	Performance Measures and Data		
IX.	Budget		
Appendix A.	Glossary		
Appendix B.	Codes		
Appendix C.	Points of Contact		
Appendix D.	Academic Alliance Centers		
Appendix E.	ASC Obligation/Cost Plan		

Contents

List of Tables

List of Figures

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 (D&E) 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 (3D) 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 (DOE), 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 DoD'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 (R&D) 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 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 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 two-dimensional (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 with Sandia National Laboratories (SNL)/Cray and Lawrence Livermore National Laboratory (LLNL)/IBM have transformed the supercomputer industry. By mid-2007, there were 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 Los Alamos National Laboratory (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 the Mesa capability computing platform, the first platform to be hosted through ACES, to be procured and sited at LANL.

In FY09, ASC released improved codes to support stockpile stewardship and other nuclear security missions, including secure transportation, NWC infrastructure, and nuclear forensics—specifically, a suite of physics-based models and high-fidelity databases were developed and implemented to support National Technical Nuclear Forensics activities.

In FY10, ASC will continue to deliver science-based simulation tools to support annual assessments and the next generation of LEPs. ASC will also provide tools for both experiment and diagnostic design to support the indirect-drive ignition experiments on the NIF and for improved confidence and response time for questions of vital importance to achieving predictive simulation capability. In addition, ASC will continue to provide national leadership in HPC and deploy capability and capacity platforms in support of Defense Programs campaigns.

By FY11 and beyond, ASC will focus on strengthening the science-basis and driving down uncertainties for weapons simulations to a degree that NNSA can ultimately, and credibly, claim predictive capability; instituting a robust, formalized peer review system; increasing the number of production computing cycles to support increased use of simulation in stockpile activities and reliance on UQ in weapons decisions; and pursuing exascale computing to meet time-urgent, future capability needs as documented in the *ASC Roadmap* and the *Predictive Capability Framework* (PCF).

III. Accomplishments for FY08–FY09

ASC accomplishments from Quarter 4, fiscal year 2008, through quarter 3, fiscal year 2009, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

ASC headquarters (HQ) is pleased to highlight the outstanding achievements of the Defense Programs laboratories.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Approved the Sequoia Critical Decision (CD) packages, CD-2 and CD-3, on January 8, 2009, with final contract award to IBM on January 12, 2009. The Sequoia Initial Delivery (ID) system (Dawn) was accepted on March 27, 2009, and began running initial tri-lab science applications on April 1, 2009. The Sequoia scalable application preparation project was initiated.
- Completed Level 2 milestone to integrate, test, accept, and deploy the Sequoia ID (Dawn) platform for initial science runs. This work included integrating new hardware and software technologies to provide the network and file system.
- Deployed job scheduling (Moab) and resource management (SLURM) services on the Sequoia ID (Dawn) (also part of above Level 2 milestone), to allow batch jobs on the Sequoia ID (Dawn) to be submitted, scheduled, and run through Moab/SLURM. Also completed the modifications to Moab and SLURM to support the Workload Characterization Level 2 milestone.
- Hosted Risk Management Techniques and Practices Workshop for High Performance Computer Centers. The workshop, jointly sponsored by NNSA and DOE Office of Science, focused on the risk management aspects of acquisition of a leading-edge computing system by bringing together managerial and technical project staff from major U.S. computing facilities.
- Deployed and began using Tripod Operating System Software (TOSS) 1.2 and OFED 1.3 at the tri-labs.

LANL Accomplishments for Computational Systems and Software Environment

- Completed the Roadrunner Level 2 milestone for final system delivery, acceptance, and science runs. This milestone brings compute power for the ASC program into the petaFLOPS range, enabling scientific calculations in support of the ASC Roadmap. Additionally, it provides compute cycles and speed to weapon's simulation and physics codes for taking the predictive capability to the next level.
- Completed CD documents, CD-0 and CD-1, for the New Mexico ACES Mesa capability computing platform. Received signature and approval for Mesa CD-1 by NNSA Principal Assistant Deputy Administrator for Military Applications Brigadier General Garrett Harencak on June 16, 2009. Mesa is the next ASC capability platform that will replace functionality now provided by the ASC Purple and Red Storm platforms.

- Deployed acceptance, stabilization, and production of Hurricane, Turing, and Lobo capacity clusters. TOSS stack and tools deployed on clusters.
- Designed, developed, and initially deployed the Parallel Log Structured File System (PLFS). PLFS can improve checkpoint bandwidth for any large parallel application that writes to a single file. The expected improvement is especially large for applications doing unaligned or random input/output (I/O). Initial results for LANL codes are showing a 2–28 times improvement on moderate sized runs on one connected unit. Additional gains are expected on larger sized runs (multiple connected units).

SNL Accomplishments for Computational Systems and Software Environment

- The New Mexico ACES received NNSA approval for the Mesa's Mission Need and Preliminary Baseline Packages, CD-0 and CD-1. This achievement allowed ACES to move forward with the formal request for proposal (RFP) to prospective vendors.
- Quantified the actual throughput gain of NNSA applications using the upgraded Red Storm quad-core processors. A code run on quad core processors takes an average of 1.113 times longer than running on an equivalent number of dual core processors. This is due to increased contention for memory and network resources. However, since 6240 dual core sockets were replaced with quad core sockets, this represents an increase of 38.4 percent in the capacity of the machine for an 8 percent increase in machine investment. Performance characterization of Red Storm quad-core upgrade was completed in FY09Q1.
- Delivered software required by Level 2 milestone 3160, including crater analysis capability that complements Level 2 milestone results from FY09. This leverages the scalable scientific and information visualization capabilities within ParaView to enable analysts to have a broader understanding of ensembles of runs by providing additional quantitative data from a simulation run, including abstract features like fragments and craters. This general capability is being applied across several codes at SNL.
- Discovered precursor anomalous behavior that correlates well with a major Tri-Lab Linux Capacity Clusters (TLCC) failure mode (out of memory). The discovery was made using the OVIS software and system monitoring tools after deploying a hardware, resource, and operating system metrics collection system on the TLCC platform Glory. This discovery represents a significant step toward the ability to predict failures on HPC systems.

Tri-Lab Accomplishments for Computational Systems and Software Environment

- Successfully completed the Q4 FY08 Tri-Lab Common Computing Environment (CCE) Multi-Site Target
- Released TOSS 1.2 (Based on Red Hat Enterprise Linux (RHEL) 5.3) production version for general availability (GA)
- Delivered Open|SpeedShop (O|SS) with new features and increased overall stability and robustness; deployed (O|SS) on most tri-lab TLCC platforms where it has demonstrated the capability to routinely analyze application runs on more than one thousand processors using its offline mode
- Deployed initial version of Workload Characterization (WC) Tool at all three labs, collecting platform usage and capturing demand, and generating reports of platform usage and demand

- Deployed MonitorJob (MoJo) combined Application Monitoring tool on TLCC machines at SNL; deployed initial version at LANL; performed testing on Roadrunner and Cerillos ASC platforms at LANL and developed preliminary list of additional data to be monitored in subsequent versions of MoJo
- Fully integrated the Gazebo Test and Analysis Suite into the TOSS stack so that Gazebo is now part of each subsequent release of TOSS

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- Moved 40 racks of BlueGene/L to the open side, creating an unclassified resource of greater than 200 teraFLOPS, which is critical for unclassified science.
- Deployed commercial identity management solution phase I. Deployed features include account creation, modification, and deletion on both unclassified and classified systems. This system results in an order of magnitude speedup in account request processing time.
- Completed integration of the Sequoia ID (Dawn) and began self-maintenance of the Sequoia ID (Dawn). This includes spares inventory and operations staff training and documentation. Developed documentation and training materials for the Sequoia ID (Dawn).
- Completed the negotiation of the Sequoia contract. Completed the west room 7.5MW electrical distribution expansion from the electrical yard into the first level of the machine room. This is necessary for future Sequoia siting requirements.
- Retired Lilac and ALC Linux clusters; TLCC clusters provide replacement cycles. Deployed and self-maintained 18 TLCC scalable units (SU): Juno, Eos, Minos on the classified side, and Hera on unclassified side.
- Hosted a Predictive Science Panel meeting, the annual tri-lab meeting with CEA (the Commissariat à l'Énergie Atomique) and a Risk Management Techniques and Practice Workshop, jointly sponsored by NNSA and the Office of Science.

LANL Accomplishments for Facility Operations and User Support

- Completed final design for the Nicholas Metropolis Center for Modeling and Simulation (known as the SCC) Infrastructure Upgrades Project. This project will add 7.2 MW of raw power to the machine room floor and includes water cooling as an additive alternative to the bidding package. This project will be completed under a Level 2 milestone in FY10.
- Deployed and switched to Request Tracker, a Web-based ticket tracking system. Setup, deployed, and integrated with Request Tracker an instant messaging system. These changes enable real-time notifications to be sent to on-duty consultants when users request help or when problems occur on ASC HPC resources.
- Completed the Level 2 milestone for deploying High-Performance Storage System (HPSS) v7.1 in the unclassified network. The milestone calls for deployment in the classified network, which will happen by the FY09Q4. This version of HPSS has major metadata performance improvements that will enable faster archive storage processing for ASC users.
- Developed initial network plan for Mesa and its file system based on performance targets in the Mesa RFP Functional Requirements draft document.

- Successfully re-accredited all classified systems pursuant to NNSA Policy Letters, known as NAPs. An in-depth risk assessment was performed on the classified systems as part of the accreditation process. Internal procedures have been updated based on the results of the risk assessment.

SNL Accomplishments for Facility Operations and User Support

- Transitioned Red Storm to support other government agency (OGA) opportunities per NNSA/ASC guidance.
- Achieved GA of TLCC capacity systems Whitney (SNL/CA), Unity (SNL/NM), and Glory (SNL/NM.
- Increased Wide Area Network (WAN) reliability by completing fully redundant fiber paths, eliminating one single run connection that remained in California.

Academic Alliances

University of Chicago Accomplishments

- Released FLASH 3.1.1 on January 15, 2009, and FLASH 3.2 on July 1, 2009. The new capabilities in these releases include 1) an unsplit hydro-dynamics solver, 2) an ability of the equation of state (EOS) unit to operate on all supported grid data structures, and 3) an enhancement in the I/O unit to the collective mode of operation in HDF5, and 4) a new add-on unit to compute radiation transfer using hybrid characteristics.
- Completed a verification study of buoyancy-driven turbulent nuclear combustion that characterized the properties of this phenomenon, which plays a key role in SNe Ia. The results show that simulations do not need to resolve the Gibson scale (the smallest length scale at which turbulence can distort the flame surface) to accurately determine the nuclear burning rate, a result that reduces the uncertainties in the properties of SNe Ia predicted by simulations.
- Demonstrated that the amount of energy released during the deflagration phase, and therefore the amount of nickel produced (and the peak luminosity of the supernova) depends sensitively on the mass of the progenitor white dwarf star—a completely unexpected result.
- Carried out the first 2D, whole star simulations of SNe Ia with sufficient spatial resolution to observe in situ initiation of a detonation wave.
- Carried out more than 35 high-fidelity, 3D, whole-star simulations of the three current models of SN Ia: the pure deflagration, deflagration to detonation transition, and gravity confined detonation. The simulations reveal that these explosion mechanisms have different and robust signatures, in terms of the predictions they make for correlations among the observed properties of SNe Ia. This holds great promise that comparisons of simulations and high-quality observations can discriminate among these mechanisms.

University of Illinois at Urbana-Champaign Accomplishments

- Licensed *Rocstar*, which is being successfully distributed. The University of Illinois spinoff company, IllinoisRocstar LLC, now provides simulation software development, support and analysis to government agencies and U.S. industry (for example, aided NASA and U.S. rocket companies in assessing the magnitude and frequency of acoustic vibrations in Ares I/Constellation (RSRMV) solid rocket

launch vehicle). Using *Rocstar's* successful model, CSAR has been participating in several National Science Foundation- and industry-led efforts to advance the state of academic and industry strategies for multiphysics coupling and massively parallel applications. The *Rocstar* multiphysics simulation application and integration framework continues to be an essential component of several advanced simulation capabilities spanning multiple departments and disciplines at the University of Illinois (for example, *RocfloCM* and *Roctherm*).

- Studied, refined, and improved the *Rocstar* technologies required to bring full-rocket length and time-scales within reach of full-physics simulations (for example, time zooming, complex geometrical case constraints, meshing and remeshing).
- Implemented and published novel and robust methods and models for "capturing" material boundaries in shock-material simulations and propellant burning (for example, erosive burning of solid propellants, applicable to both homogeneous and heterogeneous propellants).
- Based on real experiences, CSAR code development and existing applications are evolving to support modern architectures, including hybrid message passing interface (MPI)/OpenMP strategies with an eye toward petascale systems, architectures, and applications.

University of Utah Accomplishments

- The Uintah simulation software developed by C-SAFE has been formally released as a Unix/Linux tarball (version 1.1) at http://uintah.utah.edu. The software is licensed under the MIT Open Source License. The Uintah documentation is broken up into three documents: installation guide, user guide, and developer's guide, to assist users to install the code, perform test simulations and visualizations, use the code for their own simulations, and modify the code to add new features and capabilities.
- A center-wide effort to improve code performance and scaling of Uintah was completed in preparation for running on petascale computing architectures. Strong scaling of the code to 32,000 cores for the AMR-ICE algorithm was achieved by implementing algorithmic improvements to the load balancing code module. Application to AMR-MPM-ICE is now under way. Additional performance improvements were realized by rewriting the task graph module to account for variations in communications. Detailed investigations of the MPM and ICE algorithms led to a new understanding of the convergence behavior of the MPM method and to improvements in the ICE algorithm to eliminate non-physical oscillations on some speed compressible flow problems.
- C-SAFE officially adopted the open source VisIt visualization framework from LLNL for use in Uintah simulations. C-SAFE team members actively contributed to this leading-edge software project.
- The Uintah MPM-ICE component was used as an investigative tool to gain insight into the behavior of shaped charges used by oil producers. Large numbers of these charges are used routinely in oil exploration to improve the mass flow rate of oil through the well casing. Initial proof of capability simulations were performed and the results are promising. The Uintah-MPM component was utilized to study hypervelocity particulate impact on granular geo-materials, for example, sandstone.

California Institute of Technology (Caltech) Accomplishments

- Completed two main tasks towards the Level 1 milestone *Lagrangian ballistic penetration application in the Virtual Test Facility (VTF)*: 1) implementation of an

element erosion algorithm enabling the fast simulation of ballistic penetration; and 2) implementation of a non-smooth contact algorithm enabling the contact between evolving rough surfaces.

- Completed seventy-five percent of the first integrated UQ campaign (milestone 2). This first UQ run focused on ballistic penetration of steel plates by steel spherical impactors. The verification diameters for these parameters (and additional parameters such as mesh size, time step, and yield stress) were computed using the commercial finite-element code ABAQUS (as a temporary stand-in for the VTF) and the in-house UQ pipeline, which is based on SNL's Dakota optimization suite. The UQ pipeline was deployed to large computers at LLNL (Hera) and LANL (Coyote and Lobo).
- Developed an UQ framework that exploits the modular/hierarchical system for purposes of the efficient computation of probability-of-failure upper bounds of the concentration-of-measure type.
- Developed sharper probability-of-failure upper bounds of the concentration-of-measure type for systems whose response functions are strongly nonlinear and exhibit discontinuities, for example, cliff-like behavior.
- Implemented and verified (using standard benchmark tests) an Eulerian formulation of the equations of finite deformation elasticity and viscoplasticity while working within a fully multiprocessor, load balanced, Cartesian-based, block-structured adaptive mesh refinement (AMR) framework (AMROC).

Purdue University Accomplishments

- Developed an integrated MEMOSA (Microelectromechanical Systems (MEMS) Overall Simulation Administrator) solver environment to address MEMS structural response under the action of fluid forces. Structural response was addressed using an existing solver based on the material point method, while the fluid was addressed by developing a new C/C++ solver based on the finite volume method (FVM). Integration was achieved using the immersed boundary method, and the solvers were coupled through Python using MPI. Parallelization of MEMOSA-FVM was completed.
- Developed a Phase Field Micro-Mechanics solver, MEMOSA-PFMM, and implemented it in parallel to compute dislocation dynamics, and to provide macro-scale material models for the simulation of structural deformation and creep.
- Developed a verification and validation (V&V) suite for the LAMMPS molecular dynamics code and a suite of V&V for MEMOSA codes. V&V activities were initiated and demonstrated good agreement with data and theory.
- Computed structural response of MEMS structures under fluidic damping and validated it against data. Q-factor predictions within 5 percent of data were obtained for the continuum regime, well within target error bounds. Q factors for the rarefied regime were also simulated, ahead of FY08–09 goals, with yield results within 16 percent of measurements, which is already within the 20-percent error target set for FY10–11.
- Conducted large-scale simulations of metal-metal and metal-dielectric contact simulations using LAMMPS molecular dynamics code with existing interatomic potentials and software. Results for pull-out force were found to be within 30-percent of published data, within the target set for FY08–09. Results for contact area evolution were also computed.

Stanford University Accomplishments

- Completed QMU analysis to describe operability limits of canonical scramjet based on low-fidelity models. Submitted the analysis results to the AIAA Journal of Propulsion and Power to solicit critical feedback from other scientific communities.
- Completed version 1.0 of the C++ infrastructure designed to support the Center's unstructured solvers Joe, Ray, and Charles. Joe is an unstructured Reynolds-Averaged Navier Stokes (RANS) solver for full system simulation; Ray provides embedded support for UQ; Charles is the high-fidelity solver for large-eddy simulation (LES).
- Developed novel UQ propagation methods addressing two fundamental problems: analysis of high-dimensional systems, and the construction of irregular—and possibly discontinuous—response surfaces.
- Verified LES using explicit filtering demonstrated for a turbulent channel flow using a structured grid solver.
- Developed a novel, unstructured discretization for compressible flow LES that has low dissipation suitable for resolved turbulence and can accommodate arbitrary unstructured meshes and mesh transitions by using the summation-by-parts operators and associated stability concepts developed under the previous program.

University of Michigan Accomplishments

- Applied UQ software from LLNL, LANL, and TAMU to the one-dimensional (1D) Center for Radiative Shock Hydrodynamics (CRASH) problem.
- Modified the BATSRUS hydrodynamics to: 1) solve hydrodynamic equations by incorporating equation of state and ionization models, 2) use a level-set algorithm to track material interfaces, 3) incorporate a radiation transport calculation using single-group diffusion, and 4) solve axisymmetric (r,z) problems.
- Performed a nominally identical set of experiments and measurements throughout an entire experimental day on the Omega laser, to assess experimental variability.
- Developed in the PDT radiation transport code: 1) a multi-frequency transport solver working with tables, analytic functions, or fits for opacities and specific heats; 2) various options for first-order and second-order time differencing; 3) various options for appropriately time-centering opacities and specific heats; and 4) the software to enable PDT to be called from BATSRUS.
- Implemented numerous tests of software functionality. Performed initial sensitivity studies to help define the statistical approach to working with large numbers of interacting variables.

University of Texas, at Austin, Accomplishments

- Implemented and verified code across multiple modeling domains: 1) ablation modeling (CMA-type locally 1D), 2) hypersonic flow code (Initial verification of DPLR), and 3) radiation.
- Completed loose two-way coupling of radiation and ablation to hypersonic flow code (DPLR).
- Completed initial sensitivity analysis of ablation rate and peak heat flux at the full system simulation level.

- Developed and implemented delayed rejection adaptive Markov chain Monte Carlo sampling algorithms for Bayesian inversion into the Predictive Engineering and Computational Sciences (PECOS) toolkit for UQ (QUESO).

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 CCE. Achieving this deployment objective requires sustained investment in applied R&D 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 FY10

- Production capability cycles per Capability Computing Campaign (CCC) guidance
- Mesa capability computing platform delivery
- ASC Level 2 milestone Mesa capability computing platform integration readiness completion
- Full production availability of Red Storm, through the end of its existing maintenance contract (11 August 2010)
- Integrated technology trends in the design of upcoming ASC capability systems
- Evaluation of vendor offerings for compatibility with ASC capability computing requirements
- Initial technology development projects in support of future ASC capability platforms

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 CCC capabilities to the ASC customers. Projected system lifetime is five years to the end of calendar 2011. This project is currently in production operational mode and receives funds for ongoing IBM maintenance.

In FY09, LLNL delivered capability computing cycles to the tri-lab ASC scientists for the selected CCC-4 and CCC-5 projects.

Planned activities in FY10:

- Continue Purple system maintenance as required

Expected deliverables in FY10:

- Production capability cycles as per CCC guidance

Preliminary planned activities in FY11:

- Continue Purple system maintenance as required; status of the Mesa capability platform will inform programmatic decisions about the length of time Purple remains in production and under maintenance in FY11

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 current and advanced systems and strategic planning for supporting infrastructure. The major focus is to define requirements and potential system architectures for advanced systems platforms that meet ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts. In FY10, this project will focus on the project management of the Roadrunner system.

In FY09, LANL delivered the Roadrunner Phase 3 system, completed acceptance testing, and integrated the machine into the network environment.

Planned activities in FY10:

- Complete transition to operations for the Roadrunner platform
- Complete CD4b for Roadrunner
- Provide system operation in support of the weapons program

Expected deliverables in FY10:

- ASC Level 2 milestone for operational readiness for Roadrunner

Preliminary planned activities in FY11:

- Provide system operation in support of the weapons program

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 284 teraFLOPS of peak processing capability. The machine uses 2.4 GHz dual-core AMD Opteron processors and 2.2 GHz quad-core AMD Opteron processors

and a custom, very high performance, 3D-mesh communication network. Red Storm has a total of 13,600 nodes (mixture of dual and quad core) and over 75 terabytes of memory and approximately 2.35 petabytes of high-performance local disk split between two separate classification environments. Red Storm produced a 204.2 teraFLOPS result on the November 2008 HPL benchmark, which is roughly 72 percent of its theoretical peak performance. Specifications and an expanded description are available at http://www.sandia.gov/ASC/redstorm. There are now 33 sites and 61 systems based on the Red Storm architecture. Red Storm continues to support classified computations for the nuclear weapons program but is also serving as the prototype system for exploring support for OGAs at the behest of the NNSA/ASC program office.

In FY09, SNL provided system modifications to support the external file system concept for Red Storm (Dark Storm). This permits multiple need-to-know groups the ability to segregate data by physical replication of Lustre File System servers while having access to the compute capabilities of Red Storm (Dark Storm). SNL also validated full scalability of the upgraded Red Storm, scaling applications out to 38,000 processors, significantly more than the previously available 25,000.

Planned activities in FY10:

- Renegotiate maintenance and support contract with Cray, Inc.

Expected deliverables in FY10:

- Initial OGA capability

Preliminary planned activities in FY11:

- Continue production computing and complete transfer of resource to OGA projects

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

The Mesa capability computing platform is being acquired and deployed under the ACES. ACES is a joint collaboration between LANL and SNL defined under a Memorandum of Understanding to provide a user facility for capability computing to the NNSA weapons programs in support of stockpile stewardship, to develop 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 has developed tri-lab requirements and architectural specification for Mesa. The architecture and design of Mesa will be optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs.

The budget dollars in this element include platform dollars associated with payments to the selected vendor for the system contract.

This project covers all aspects of program and procurement planning for current and future capability systems and strategic planning for supporting infrastructure. A major effort is to continue to identify tri-lab capability requirements.

In FY09, LANL and SNL identified the Mesa system architectural requirements and used the NNSA Chief Information Office project execution model for procurement of Mesa system.

Planned activities in FY10:

- Use the NNSA Chief Information Office project execution model for the acquisition and integration of Mesa
- Complete CD-2 and CD-3 for Mesa
- Complete ASC Level 2 milestone for system integration readiness for the Mesa platform
- Complete system integration of Mesa

Expected deliverables in FY10:

- Contract award for Mesa
- Delivery of Mesa
- Mesa ASC Level 2 system integration readiness milestone
- System integration of Mesa into the LANL computing environment

Preliminary planned activities in FY11:

- Complete Mesa acceptance testing, including the demonstration of key NNSA applications performance as defined in the contract
- Transition Mesa to operations
- Complete CD-4 for Mesa
- Complete Mesa tri-lab ASC Level 2 milestone for production capability readiness

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 Mesa RFP, equivalent analysis is required 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, and support new FastForward technology development projects to address these issues. The FastForward projects are targeted to start in FY12 or possibly FY13.

In FY09, LANL and SNL completed specification for next capability system to address anticipated workload requirements, and defined and establish a D&E project with the Sequoia runner-up that is focused on quantifying the application impact of MPI enhancements in successive generations of advanced interconnection network technology.

Planned activities in FY10:

- 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
- Provide technical oversight for the interconnection network D&E project

Expected deliverables in FY10:

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

Preliminary planned activities in FY11:

- Begin analysis for follow-on capability system
- Develop exascale roadmap

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 FY10

- Initial tri-lab TLCC capacity procurement preparation, including requirements gathering and architecture requirements, for a possible FY11 TLCC procurement
- Maintenance and operation of TLCC07 systems in production computing environment running the TOSS stack
  - At LLNL: Juno (a cluster of 8 SU), Eos (2 SUs), Hera (6 SUs, 4 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)
- Integrated CCE selected stack and tools on TLCC clusters

WBS 1.5.4.2-LLNL-002 Capacity Computing Planning and Integration

The LLNL ASC strategy for capacity computing is to leverage industry advances and open source software standards to build, field, and integrate Linux clusters of various sizes into classified and unclassified production service. 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 load of ASC scientific simulation capacity workloads. In FY10, LLNL will continue to monitor the computing industry developments for opportunities to augment capacity computing and the associated infrastructure.

Planned activities in FY10:

- Execute a follow-on TLCC capacity procurement in FY11–FY12, if directed to do so by HQ; this procurement would be handled by the TLCC technical team

Expected deliverables in FY10:

- Initial TLCC capacity procurement preparation if there is to be a FY11 capacity procurement

Preliminary planned activities in FY11:

- Continuation of FY10 activities

WBS 1.5.4.2-LANL-001 Capacity System Integration

The Capacity System Integration project is responsible for integrating new ASC capacity systems into LANL classified, unclassified, and open computing networks. Included in this project is support for the ASC capacity system acquisition strategy along with providing requirements that help to achieve the strategy. In FY10, the main focus will be on production support of capacity systems at LANL and planning for future TLCC architectures and integration.

In FY09, LANL accepted and deployed Hurricane, Turing, and Lobo TLCC clusters; the TOSS stack and tools were deployed on these machines.

Planned activities in FY10:

- Provide production support of LANL TLCC systems (Hurricane, Lobo)
- Gather tri-lab requirements and architecture requirements for FY11 TLCC procurement

Expected deliverables in FY10:

- Production support of LANL TLCC systems (Hurricane, Lobo)
- Tri-lab requirements gathering and architecture requirements for FY11 TLCC procurement

Preliminary planned activities in FY11:

- Deliver and integrate TLCC capacity computing systems

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.

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.

In FY09, SNL deployed three TLCC common capacity clusters into production with a common operating system and computing environment.

Planned activities in FY10:

- Continue maintenance and operation of TLCC07 systems in SNL's production computing environment running the TOSS stack: Glory (2 SU), Whitney (2 SU), and Unity (2 SU)
- Integrate CCE software stack and tools deployed on TLCC clusters
- Support TLCC11 procurement process for the next generation of ASC capacity computing systems

Expected deliverables in FY10:

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

Preliminary planned activities in FY11:

- Continue maintenance and operation of TLCC07 systems in SNL's production computing environment running the TOSS stack
- Deploy TLCC11, the next generation of ASC capacity computing systems, if available

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 FY10

- BlueGene/Q R&D prototype hardware release
- Evaluation of Sequoia benchmark suite on simulated environment
- Early Sequoia software functionality delivered for installation on Sequoia ID (Dawn)
- Sequoia ID (Dawn) system transition to the secure computing facility (SCF)
- Installation and deployment of new Hyperion storage and networking hardware
- Transfer of the Roadrunner system to operations with the completion of the CD-4 process; the system is running with a select set of users; complete infrastructure support provided including: I/O, networking, and archiving; the advanced Architecture operation state software stack is provided; the system is running qualified hybrid applications; and ASC and weapons science applications target for Roadrunner have demonstrated operational readiness
- Work with industry, academia, and laboratory partners to develop advanced memory and high-speed interconnect subsystems to improve performance on legacy and new ASC applications, explore revolutionary new system software approaches

for extreme scale systems, and engage in or hold workshops to achieve a common vision with other HPC sites

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

The BlueGene/P and BlueGene/Q R&D 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 IBM Engineering and Technology Services Division, Argonne National Laboratory, and the ASC tri-labs. In FY08, a BlueGene/P system was delivered to Argonne. In FY09, a BlueGene/P system was delivered to LLNL as the first system (the Sequoia ID (Dawn)) in the Sequoia procurement. 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.

In FY09, LLNL led the ANL/LLNL partnership in its interactions with IBM on the topic of BlueGene/Q. In addition to installing a BlueGene/P system (Dawn) at LLNL, the technical reviews performed resulted in freezing the compute node and link architectures, operating system features and tool APIs, hardware packing/cooling, and I/O architecture.

Planned activities in FY10:

- Continue technical interaction with IBM on remaining hardware issues and software development
- Investigate compiler capabilities for utilization of floating point, thread level speculation and transactional memory
- Utilize software and hardware simulation capabilities to report early estimates of benchmark code performance

Expected deliverables in FY10:

- Internal performance predictions for thread level speculation and transactional memory enhancements to benchmark codes

Preliminary planned activities in FY11:

- Perform work associated with preparing second version of BlueGene/Q chip hardware and price projections for build Go/NoGo decision

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.

In FY09, LLNL evaluated vendor submissions for the Sequoia RFP with respect to expected performance for Sequoia benchmark applications. LLNL also initiated investigations into threading performance issues with a particular focus on novel hardware features anticipated in Sequoia. Exploration of software transactional memory suggested that hardware support for this programming paradigm could benefit several ASC applications. Performance improvements were realized on current architectures for key ASC applications, with a particular focus on integrated codes.

Planned activities in FY10:

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

Expected deliverables in FY10:

- Evaluation of Sequoia benchmark suite on simulated environment

Preliminary planned activities in FY11:

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

WBS 1.5.4.3-LLNL-003 Sequoia

The Sequoia project will deploy a multi-petaFLOPS computer in early FY2012 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 12–24 times more capable than ASC Purple for UQ and up to 20–50 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 activity awarded build, delivery, development, and engineering contracts to IBM on January 8, 2009.

There are two major deliverables for Sequoia. The first deliverable (acquisition and delivery of an early environment—called the Sequoia ID (Dawn)) was accepted in CY2009Q1 and dedicated on May 27, 2009. The second deliverable (acquisition and delivery of final Sequoia environment) is to be completed by end of 2011. These environments (both the Sequoia ID (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 the Sequoia ID (Dawn) and Sequoia, including

visualization engines, archival storage movers, BlueGene/L, Purple, and TLCC07 Linux clusters.

Supporting the final Sequoia platform deliverable is a D&E contract for hardware, software, and scalability items. The Sequoia acquisition RFP process requested technology roadmaps from responding vendors, with the intent that the Sequoia project would fund D&E efforts for the winning vendor to address the identifiable risks and issues for platform build and delivery. The D&E contract with IBM is intrinsic to overall Sequoia risk mitigation strategies and helps to ensure the successful "productization" of the Sequoia components and systems.

Planned activities in FY10:

- Manage FY10 D&E contractual milestones
- Hold review(s) for Sequoia planning
- Plan for Sequoia platform delivery
- Prepare for Sequoia build Go/NoGo

Expected deliverables in FY10:

- Early Sequoia software functionality
- Sequoia ID (Dawn) system transition to the SCF
- Demonstration of first and second pass Sequoia prototypes

Preliminary planned activities in FY11:

- Manage FY11 D&E contractual milestones
- Plan for Sequoia parts commit and options
- Plan for Sequoia demo and early science runs

WBS 1.5.4.3-LLNL-004 Hyperion Test Bed

With the extremely demanding IO requirements of petascale applications for Sequoia and the need for improved scientific data management capabilities, it is clearly apparent that emerging breakthrough technologies such as storage server virtualization and flash memory need to be tested in a large-scale environment such as Hyperion. The Hyperion Test Bed project will work with an expanded set of Hyperion vendor partners in the Hyperion Data Intensive Test Bed to evaluate innovative alternative storage and storage area network architectures and push the achievable bandwidth input/output operations per second performance boundaries. In addition, it will work with the Hyperion Data Intensive Test bed to improve the fail-over and reliability of parallel file systems.

Planned activities in FY10:

- Procure storage and networking hardware for the Hyperion Data Intensive Test Bed

Expected deliverables in FY10:

- Installation and deployment of newly acquired storage and networking hardware

Preliminary planned activities in FY11:

- Procure storage and networking hardware for the Hyperion Data Intensive Test Bed

WBS 1.5.4.3-LANL-002 Roadrunner Phase 3 Support

The budget dollars in the Roadrunner Phase 3 Support project are platform dollars that are associated with payments to IBM for the system contract or the analyst support. 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.

In FY09, the Roadrunner Phase 3 system was integrated into the LANL network environment, which included the completion of science runs as part of system stabilization efforts. The Roadrunner Phase 3 system transitioned into the classified network.

Planned activities in FY10:

- Provide system and application analyst support for Roadrunner

Expected deliverables in FY10:

- Transition of Roadrunner Phase 3 to operational status

Preliminary planned activities in FY11:

- Provide system and application analyst support for Roadrunner

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

The Advanced Systems Technology R&D project will work in conjunction with the 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. The project's architecture efforts will focus 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 also explore software capabilities that will be essential for increasing the performance and scalability of applications beyond those provided by general-purpose operating systems, while providing the functionality necessary to deal with new compute node architectures, networks, parallel programming models, and applications—all at extreme scale.

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.

In FY09, SNL added several new capabilities. To support alternative programming models and many-core processors, the project created an application programming interface usable by software libraries to implement shared memory in lightweight kernel (LWK) operating systems. The new application programming interface means that hardware-specific capabilities can be abstracted for use by the developer. Additionally, a lightweight debugging mechanism for extreme-scale applications has been developed and deployed. It focuses on debugging a currently running, but hung, job. SNL also established a strategic collaboration among DoD, Micron Technology, and the University of Maryland; to use Structural Simulation Toolkit (SST) and DRAMsim to analyze advanced memory concepts and quantify impact on ASC applications. Finally,

SNL developed a new capability, using field programmable gate arrays (FPGAs) to study advanced memory controller designs and in particular the impact on sparse matrix-vector operations, and deployed a new open source micro-benchmark that makes realistic measurements of a high-speed interconnects messaging rate capabilities

Planned activities in FY10:

- Work with industry, academia, and government laboratory partners to develop advanced-memory and high-speed interconnect subsystems to improve performance on legacy and new ASC applications
- Explore revolutionary new system software approaches for extreme scale systems
- Engage in or hold workshops to achieve a common vision with other HPC sites

Expected deliverables in FY10:

- New research directions based on FY09 IAA research plan
- Specification and implementation of an advanced memory subsystem using the SST simulation environment
- An open-source messaging rate micro-benchmark deployed that can be used by the research community to measure realistic message rates for advanced high speed interconnects
- Demonstration of the acceleration of a key solver operation, sparse matrix-vector multiplication, using an advanced memory controller design in an FPGA test bed

Preliminary planned activities in FY11:

- Identify and develop new research directions, in cooperation with the IAA, to impact future capability platforms

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 FY10

- Testing and support of Open Fabric Enterprise Distribution 1.4, part of RHEL 5.4
- Deployment of TOSS 1.3 (based on RHEL 5.4)
- Improved Sequoia ID (Dawn) system applications development tools environment
- Enhanced production version of highly scalable code correctness tool suite
- Certification of the TLCC Tripod applications development tools environment

- Design and prototype of a domain-specific language (DSL) focused on a specific set of low-level operations used by a particular method/domain (for example, hydro) within the ASC program; the DSL will include a cost model that will support heuristics for scheduling regions of the program on heterogeneous computing resources
- Report on the performance analysis of Mesa
- Initial study on the performance of possible exascale systems
- Implementation of the cooperative caching file system and documentation of results from testing the cooperative caching file with benchmarks and real applications
- A document that specifies the Joint Exascale Initiative (JEI) project definition, outlines future work, milestones, and operating principles; interaction with the broader community in the specification of exascale R&D activities
- A successful execution of test applications as preparation to transition Roadrunner Phase 3 system to an operational state after it has moved to the secure and to accept Mesa into production
- Engagement with the Sequoia and Mesa teams as well as the IAA to provide the capability to simulate advanced architectures to help mitigate the risks in moving to new platforms. This will include extensive validation efforts for all components, as well as continuing and expanding community outreach efforts.
- System software developed as risk mitigation should the vendor-provided Mesa software fail to perform
- Application scalability performance studies as part of the Mesa acceptance criteria to support the Mesa capability computing platform acquisition of Mesa
- Support for evolutionary system software for current HPC hardware

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. LLNL uses TOSS on all of its Linux clusters, not just TLCC. This project also funds approximately 60 percent of the manpower required to develop, deploy, and maintain TOSS. The funding LLNL receives for its portion of CCE TOSS funding accounts for 40 percent of the LLNL effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in section 1.5.4.7 of this document.

In FY09, LLNL deployed job scheduling (Moab) and resource management (SLURM) services on the Sequoia ID (Dawn) (also part of above Level 2 milestone), to allow batch

jobs on the Sequoia ID (Dawn) to be submitted, scheduled, and run through Moab/SLURM. LLNL also completed the modifications to Moab and SLURM to support the Workload Characterization Level 2 milestone. LLNL employed and began using TOSS 1.2 and OFED 1.3 at the tri-labs.

Planned activities in FY10:

- 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 FY10:

- Deployment of TOSS 1.3 (Based on RHEL 5.4)
- Deployment of OpenFabric Enterprise Distribution 1.4
- Initial work to port Moab and SLURM to Sequoia ID (Dawn) system

Preliminary planned activities in FY11:

- Continue ongoing TOSS software development and support
- Deploy TOSS 2.0 (based on RHEL 6)
- Develop InfiniBand storage area network (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 (ADEPT) 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.

With the deployment of the Sequoia ID (Dawn) system applications development tools environment in FY09, the ADEPT project is evaluating user and code developer needs in preparation for the eventual Sequoia system. This effort includes a high priority focused

activity within ADEPT, entitled Scalable Applications Preparations (SAP). This activity will ensure that a wide range of ASC applications run effectively on Sequoia when it arrives, including exploiting the novel aspects of its hardware and software. It also includes ADEPT responsibility for the SAP Level 2 milestone.

Planned activities in FY10:

- Provide, maintain, and refine Purple, the Sequoia ID (Dawn), BlueGene/L, and TLCC code development environments; Purple environment will also be maintained for the duration of its useful life (expected to be through the end of FY10)
- Coordinate integrated design code scaling for Sequoia as part of SAP activities
- Explore nested node concurrency programming model interfaces and performance as part of SAP activities
- Work with tri-lab partners to advance the common tri-lab environment for TLCC
- 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 FY10:

- The FY10 SAP milestone, including characterizing the performance on the Sequoia ID (Dawn) and anticipating Sequoia performance for one or more multi-physics codes
- Sequoia ID (Dawn) system applications development tools environment improvements
- Production version enhancements of highly scalable code correctness tool suite
- TLCC Tripod applications development tools environment certification for new TOSS releases

Preliminary planned activities in FY11:

- 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
- Continue to explore nested node concurrency programming model interfaces and performance
- Continue to support users and to interact with vendors to serve user needs
- Prepare for validation and acceptance of the Sequoia system

WBS 1.5.4.4-LANL-003 Code Performance and Throughput

The Code Performance and Throughput project is coordinated with the ASC integrated code development. The goals of the project are to measure, understand, predict, and improve the performance of the ASC integrated weapons physics code projects at each major release, and to inform and assist ongoing code development and refactorization efforts from a computer-science perspective.

The project provides capabilities for ASC integrated physics codes in three areas: 1) metrics for and prediction of ASC integrated physics code performance; 2) optimization of integrated codes; and 3) computer science input on software architecture and re-factorization decisions. To ensure relevance, this project will deploy staff directly on code projects to assist in enhancing code performance in each major release.

As an overarching goal, the project will seek to estimate the runtime and throughput based on large-scale variables such as physics method, computer system, and data sets, thus enabling the definition of targets for enhanced performance as a function of time.

In FY09, LANL's automated performance-metric-tracking tools matured and were used extensively by the Eulerian Codes project. Performance metrics were used widely by developers for code-development decisions and by management for reporting and auditing. Optimization work for the Eulerian Codes project was on-going with contributions including groundwork for the parallelization of a recently integrated code base. Simplified AMR data structures and algorithms were proposed and implemented for the Eulerian Codes project.

Planned activities in FY10:

- Enhance performance metric-tracking software on tri-lab systems, expanding coverage of codes and problems, and incorporating deeper levels of application knowledge (for example, time-step control and AMR parameters in physics packages)
- Improve Web-based interfaces to performance data for X Division (weapons codes) and CCS Division (transport) code projects
- Deploy initial performance-tracking support for hybrid computer systems (Roadrunner) and develop this further to address issues of emerging hardware and asynchronous data movement
- Deploy optimization expertise as needed to ASC integrated code teams, with an initial emphasis on supporting the Eulerian re-factorization effort
- 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
- Work with ASC physics code teams to inform, from a computer-science perspective, their on-going development and re-factorization efforts for current and future ASC platforms, and to investigate and facilitate cross-fertilization with ASC computer science projects (exascale programming paradigms, analytic performance modeling)
- Support the FY10 Level 2 milestone "Application enablement on next-generation platforms," in coordination with the Advanced Architecture and Usable Supercomputer project

Expected deliverables in FY10:

- Improved understanding of performance trade-offs in X Division code projects, and enhanced performance of FY10 releases in the Eulerian Codes project
- Improved metrics for Roadrunner and hybrid computer systems
- More physics-based and algorithm-aware metrics (for example, time-step control)
- Automated performance tracking extended to more codes, and more end-to-end problems

- On-going support for code re-factorization efforts in the Eulerian Codes project: integration of a new multi-physics code base (parallelization, and code reuse); implementations of new, more efficient data structures and algorithms in the existing code base, focusing on AMR and communication, with a goal of improving scalability and runtime, and reducing memory footprint for production problems; and support for the activities of other projects and milestones, including the Applications Readiness project and X Division code projects

Preliminary planned activities in FY11:

- Continue to improve metrics
- Improve performance of ASC integrated physics code projects
- Investigate the application of new algorithms and programming paradigms to ASC physics codes with a view to preparing for exascale computer systems

WBS 1.5.4.4-LANL-005 Software Support

The Software Support project works to establish a strong development and analysis tool capability for current and next generation HPC platforms, including parallel capabilities. It is focused on working with the HPC tool community and vendors to identify, plan, and integrate tools into production environments and establish a solid support structure. The project supports the incremental improvement of tools driven by the ASC 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.

In FY09, LANL deployed O|SS, which became functional on all production platforms.

Planned activities in FY10:

- Continue building an MPI support capability by engaging the community support model; focus will be on OpenMPI for development needs
- Identify requirements and approaches to address scaling issues as ASC moves toward Mesa
- Assess application MPI usage by platform architecture and recommend parameter additions or code enhancements
- Continue performance analysis support capability; focus will be on O|SS integration in application analysis and new analysis needs
- Increase memory analysis and MPI performance tool support
- Build a stronger debugger support capability; work with TotalView on tool deployments and enhancements, and work with LLNL debugger capability through the static analysis tool and subset debugging
- Initiate a capability to assess compiler support integration in software development strategy for users
- Continue involvement with DOE NNSA and Office of Science laboratories in petascale tool planning and development programs

- Continue efforts in LANL architecture planning efforts and Mesa deployment

Expected deliverables in FY10:

- Application assessment for MPI usage to increase throughput
- Interaction with community and vendors to increase MPI scale capability
- Direct support to code projects for performance analysis integration
- Direct support to code projects for debugger capability assessment and use
- Requirements planning for compiler efforts to support code project needs and architecture expectations
- Software Support and Development capability strategy plan updates

Preliminary planned activities in FY11:

- Continue building an MPI support capability by engaging the community support model; focus will be on OpenMPI for development needs
- Continue performance analysis support capability
- Provide memory analysis and MPI performance tool support
- Continue work with TotalView on tool deployments and enhancements, and work with LLNL debugger capability through the static analysis tool and subset debugging
- Continue involvement with DOE NNSA and Office of Science laboratories in exascale tool planning
- Continue efforts in LANL architecture planning efforts and Mesa operation

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 production problem solving (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, system configuration verification and software stack deployment with real user applications (accomplishing science during stand up), and metrics and analysis/profiling.

In FY09, LANL completed runs to enhance production maturity of Roadrunner Base system, completed runs to accept TLCC systems into production, and completed runs to transition Roadrunner Phase 3 system to an operational state in the open.

Planned activities in FY10:

- Take on the hardest and most elusive problems, or newly identified problems, with existing production machines
- Ensure an orderly application deployment on the Mesa system

- Continue to assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner Phase 3 system
- Implement special projects, including library code port to Roadrunner, dirty page checkpointing, and PLFS testing
- Assist system management personnel with problem investigation and resolution
- Continue to assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner Phase 3 system
- Implement special projects, including library code port to Roadrunner, dirty page checkpointing, and PLFS testing
- Assist system management personnel with problem investigation and resolution

Expected deliverables in FY10:

- Runs completed to transition Roadrunner Phase 3 system to an operational state after it has moved to the secure
- Runs completed to accept Mesa into production

Preliminary planned activities in FY11:

- Continue to take on the hardest and most elusive problems, or newly identified problems, with existing production machines
- 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 their productivity, capabilities, and performance on current and future ASC machines. This project proposes, implements, and tests new approaches to improve the productivity of ASC code projects on current and future computer architectures. These approaches include resilience through optimization procedure and scheduling, new mesh data structures, new strategies for AMR, performance improvement, data locality, compressed data structures for materials, and parallel strategies on advanced computer architectures.

This project provides a bridge between future computer architectures and ASC code projects for efficient usage of current and future computer hardware. The staff members in this project are able to identify main bottlenecks of current ASC code projects for better performances in calculation speed, memory usage, and communications on current and future machines, to develop strategies and plans for remedies, and to implement changes within ASC code projects. The project currently focuses on mesh data structures, material data structures, users' data structures for I/O, and AMR.

In FY09, LANL added many new capabilities, including developed a library, AMTTI, for application mean time to interrupt and machine subsystem fail rates; applied for and received a government use license for the library, and released the library to the National Security Agency (NSA) Center of Exceptional Computing; implemented and tested a framework for testing data structures within Crestone project to reduce memory footprints; developed a procedure to improve the memory usage within Crestone project; implemented a flattened mesh data structure within Crestone project; released a

more efficient I/O package, high-performance I/O, in Crestone project; demonstrated hydro calculations on cells of Roadrunner, and on multiple connected units of Roadrunner; developed a new AMR infrastructure of AMR for much better data locality; and developed a new approach for hydro in the Refactor branch within Crestone project.

Planned activities in FY10:

- Develop the flattened mesh data structures within Crestone project
- Demonstrate the advantages of flattened mesh data structures over the existing hierarchy structures within Crestone project
- Develop the weight solution scheme with the library AMTTI
- Collect application data and validate models/tools against real data

Expected deliverables in FY10:

- The library AMTTI with weight solution schemes
- The compressed material data structure within Crestone project
- The hydro capability on compressed material data within Crestone project
- Implementation of both 2D and 3D hydro capability within the new framework of AMR

Preliminary planned activities in FY11:

- Investigate additional optimization procedures on petaFLOPS machines
- Develop tools to optimize production runs on petaFLOPS machines, in coordination with the JEI
- Improve the performance of Crestone through flattened mesh data structures
- Apply compressed material data structures to other physics packages within Crestone project
- Develop hydro capability with compressed data structures within the new framework of AMR

WBS 1.5.4.4-LANL-008 Advanced Architectures and Usable Supercomputing

The scope of the Advanced Architectures and Usable Supercomputing project (the ASC component of LANL's Center for Advanced Architectures and Usable Supercomputing) is to measurably improve the usability, performance, reliability, efficiency, and productivity of petaFLOPS supercomputers for nuclear weapons applications. The goal is to guide the architectural choices and software environment for next-generation capability machines, with relevance to ASC maintained by working in the context of a radiation-hydrodynamics workload. Integrated with the JEI project (WBS 1.5.4.4-LANL-006), the Code Performance and Throughput project (WBS 1.5.4.4-LANL-001), and the integrated codes projects, this work supports ongoing Roadrunner weapons science and algorithms development.

Several important new developments in FY09 include the implementation of cell-messaging layer on Roadrunner, and load-balancing techniques for overlapping computations between Opterons and cells in Roadrunner. Large-scale system analysis included the comparison of Roadrunner, the Sequoia ID (Dawn), and TLCC clusters.

Planned activities in FY10:

- DSL tailored for ASC subject areas, and explore techniques for translating high-level, abstract code representations into low-level code for current and emerging processor designs
- Explore the design and implementation of an extended SIMD programming model for coprocessors with high-bandwidth serial interconnects, and continue the development of a novel vector computing architecture, exploring the implementation of key kernels to model full ASC applications at scale
- Design and implement a cooperative caching file system using compute-node RAM as local storage, evaluating performance under synthetic and real applications characteristic of LANL workloads
- Analyze and model the performance of expected large-scale systems deployments (including Mesa), and explore memory access requirements of applications in innovative, possibly heterogeneous, memory structures
- Initiate research that couples models of performance, power, and reliability
- Investigate inter-core communication on mainstream processors using reduced-functionality MPI (extending the cell-messaging layer)
- Initiate research on coupling performance modeling with the modeling of power and reliability
- Extension of cell-messaging layer to support mainstream multi-core processors as well as the IBM Cell-BE
- Use PALs application performance modeling capability to initiate studies of possible future advanced architectures that may be possible in the exascale regime

Expected deliverables in FY10:

- A prototype DSL for an ASC domain (hydro) and associated compilation toolset
- A prototype hybrid co-processor/reconfigurable hardware platform, and implementations of relevant ASC kernels in hardware
- Report on the performance analysis of Mesa and initial studies on the performance of possible exascale systems and innovative memory structures
- Software and methodology supporting modeling of performance, power and reliability, multi-core message passing on mainstream processors, and trace-to-benchmark extensions for the coNCePTuaL toolset
- Implementation of the cooperative-caching file system, and evaluation of effectiveness of cooperative caching for LANL workloads

Preliminary planned activities in FY11:

- Enhance usability of DSL, and explore code transformations for heterogeneous processor architectures
- Develop larger-scale system, and implement high-level language for scientific users
- Continue integration of performance/power/reliability modeling, and further explore exascale systems and memory organizations

WBS 1.5.4.4-LANL-009 Joint Exascale Initiative

The JEI is a strategic collaboration with SNL to develop the tools and capabilities required for future computing platforms that operate at the exascale in computational power and data intensity. JEI focuses on the areas critical to the rational design and usability of exascale systems, including performance analysis and modeling, programming of emerging architectures, and new computing platforms.

JEI is a new initiative between LANL and SNL. LANL and SNL will hold a joint planning meeting to kick off discussions of JEI's exascale technical R&D agenda at the Los Alamos Computer Science Symposium in October of FY10.

Planned activities in FY10:

- Stand up JEI
- Define joint projects and areas of emphasis
- Define future JEI milestones

Expected deliverables in FY10:

- JEI project definition
- Joint report from LANL/SNL outlining future work, milestones, and operating principles for JEI

Preliminary planned activities in FY11:

- Coordinate project activities with the broader trans-petascale community efforts

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 of future computational systems. A major focus this year will be preparation for Mesa. An important initiative will be creating a run time environment for Mesa as risk mitigation to the vendor-supplied solution. Scalable system software has been identified as a high-risk aspect of this capability-class system procurement. The user/application requirements for Mesa are quite a challenge for the system software. The breadth of functionality will be difficult to achieve while maintaining high parallel scalability to the size of the machine. SNL will create an LWK Operating System for Mesa based on the Kitten LWK that was originally developed under SNL's LDRD program. A compatible run time system, including job launcher and application libraries, will be developed. The initial implementation will target the hardware for a system called Red Sky, located at SNL, Albuquerque. As a leading edge, high-performance cluster, the Red Sky hardware should provide a reasonable development platform since the final target is unknown.

While preparations for Mesa are a critical activity for FY10, SNL will continue to provide evolutionary software for the existing platform. No major enhancements are planned, but software issues and user productivity bottlenecks will continually be addressed. On the other end of the planning spectrum, software R&D for potential future systems will be performed. Mesa will likely be the last capability-class system to run the ASC codes with minimal application modifications. R&D efforts will focus on system software features that can facilitate a transition to new programming models dictated by hardware advances and even more parallelism. This includes continued development of the HERMES (Heterogeneous Resource Monitoring for Extreme Execution Scalability)

framework and algorithms for intelligent resource monitoring, analysis, and allocation applicable on next-generation multi-core systems. In addition, novel programming models, which take advantage of new hardware and system software capabilities, will be explored. This project will also analyze the runtime performance on current capability and capacity platforms to 1) better understand the runtime characteristics of key applications of interest, 2) provide a better understanding of the underlying architecture of the platform being tested, and 3) ensure application performance is reasonable, if not optimal, for a given system. This effort will also apply modeling and performance prediction techniques to provide a better understanding of future architectures and platform acquisitions.

In FY09, SNL provided enhanced system software for HPC capability platforms. The revised LWK supports multi-core processors dynamically. While the previous version only supported one or two cores per node, the current release allows for "N" cores per node. The release is now running on Red Storm with four cores per node. The software has also been verified on compute modules with eight cores per node. A new, shared memory mechanism was also introduced, significantly improving MPI collective operations. A three-times improvement has been seen for messages up to 32 kilobytes in size.

Planned activities in FY10:

- Develop system software as risk mitigation should the vendor-provided Mesa software fail to perform
- Support the Mesa acquisition by providing application scalability performance studies as part of the Mesa acceptance criteria
- Support evolutionary system software for current HPC hardware
- Develop multi-core framework for resource monitoring and task/resource re-direction

Expected deliverables in FY10:

- Kitten-based LWK and run time environment as risk mitigation for Mesa; the software will run on SNL's Red Sky hardware as proof of viability and performance
- Mesa acceptance tests documentation and application performance report
- HERMES implementation and demonstration of capability

Preliminary planned activities in FY11:

- Address scalability issues in Mesa

WBS 1.5.4.4-SNL-003 System Simulation and Computer Science

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 and evaluation approach that permits decisions to be informed by predictions of application performance on these machines. Such an approach can be used in architecture design, algorithm design, and platform procurement. The System Simulation and Computer Science project, in collaboration with the IAA, will continue the development of a multi-scale simulation capability within the context of the SST. The SST consists of a core set of components that enable parallel discrete-event simulation; high-fidelity networking, memory, and processor components; and coarse-grained

simulation components that capture essential elements of machine performance with low computational cost.

Future HPC systems and the applications designed to utilize them are impacted by a variety of considerations, including: scalability of applications, ease-of-programming, memory and network latencies becoming more imbalanced relative to computation rates, data corruption and its propagation, frequency of interrupts, power consumption, and overall machine cost. SST is designed to allow each of these parameters to be explored allowing the consideration of a broad space of potential architectural and algorithmic designs. The goal is for the SST components to be extended and enhanced by a community of simulator developers including academic, industrial, and government partners. An even larger community is expected to be the users of SST, including algorithm developers, architecture designers, and procurement team members.

In FY09, SNL performed a variety of activities related to the building of a HPC platform simulation community: met with GATech and AMD to discuss ongoing and new collaboration; established a strategic collaboration among DoD, Micron Technology, and U-MD to use SST and DRAMsim to analyze advanced memory concepts and quantify impact on ASC applications; held a simulator workshop; and highlighted the simulator at the SNL Computational and Information Sciences external review.

Planned activities in FY10:

- Develop and validate parameterizable advanced memory models for future systems
- Develop and integrate NIC and router models based on the Red Storm network, and expand these models to look at alternate network topologies
- Expand SST technology models to include area and cost estimation, and validate these models
- Integrate microscale and macroscale simulation components more tightly
- Develop and integrate into SST macroscale components for network and processor models
- Engage with the Sequoia and Mesa teams as well as the IAA to find ways that the SST can assist in machine procurement and algorithm design
- Begin extensive validation efforts for all simulation components
- Continue and expand community outreach with the goal of having key community simulation components become interoperable with the SST

Expected deliverables in FY10:

- Demonstration of the use of a Red Storm network model within the parallel SST/Core framework
- Use of SST to estimate the performance of several Sequoia and Mesa acceptance tests running on the architectures we anticipate to be delivered
- Demonstration of an integrated macro-/micro-scale multiscale

Preliminary planned activities in FY11:

- Continue simulator validation efforts and explore the potential for UQ
- Develop detailed advanced processor and memory models along with coarse-grained reductions of these models
- Develop tools to ease simulator use for the primary use-cases: procurement,

application design, and architecture design

- Continue community outreach with the goal of reaching a critical mass in this year

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 FY10

- Production deployment of HPSS R7.1.1 and release of R7.1.2
- Release of HPSS R7.2
- Design documents, initial coding, and coding support for HPSS R8.1
- Requirements for next HPSS upgrade
- Production deployment of Lustre 1.8.
- Deployment of data analysis NAS file system into SCF production
- Report on parallel file system risk mitigation study
- Switching infrastructure upgrade for future Sequoia network augmentation
- Preparation and testing of I/O infrastructure for Mesa
- Deployment of PLFS on Roadrunner
- Expanded network capacities in the secure for Roadrunner and Mesa
- Parallel storage interface (PSI) requirements document
- PSI 2010 release

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, Lawrence Berkeley 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, thereby increasing 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 14 PB in a single name space.

In FY09, HPSS R7.1 was deployed to LLNL production environments and included metadata performance enhancements, tape aggregation, improved intelligence in the disk cache allocation algorithm, and parallelized class-of-service changes. Deployment teams also selected, procured, and deployed an upgrade of LLNL Sun Titanium tape drives, doubling the native titanium media capacity from 500 GB to 1 TB per cartridge.

Planned activities in FY10:

- Develop HPSS R7.1.2, a maintenance release
- Develop HPSS R7.3, a minor release including recovery from multiple hierarchy levels and advanced repack features required for repacking tape aggregates
- Develop HPSS R8.1, featuring an architecture utilizing distributed core servers and partitioned metadata to meet the extreme scalability requirements of petascale computing in the Sequoia time frame
- Develop programmatic quota system for LLNL archive users
- Deploy HPSS R7.1.1 or R7.1.2 (maintenance release) in production environments
- 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 FY10:

- Design documents and initial coding for HPSS R8.1
- Release of HPSS R7.2, featuring user defined attributes, and support for new features available with the IBM TS3500 tape library system
- Production deployment of HPSS R7.1.1 or R7.1.2 (maintenance release)
- Archival storage hardware and software procured and deployed to support expansion of ASC platforms

Preliminary planned activities in FY11:

- Develop HPSS R8.1 through code review, system test, integration test phases, and potentially beta release
- Plan for production deployment of HPSS R7.3 or R8.1
- Deploy programmatic quota system for LLNL archive users

WBS 1.5.4.5-LLNL-002 Parallel and Network File Systems

The Parallel and Network File System (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 Global Parallel File System (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, the Sequoia ID (Dawn), and Sequoia platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

The parallel and NFS project focuses closely tied development and deployment/administration teams on the task of supplying high performance, highly reliable, parallel and serial file systems in support of applications 24 hours a day. Lustre developers work hand-in-hand with the open source community to provide a scalable file system product for the ASC environment. Lustre and NAS administration teams procure, integrate, deploy, and operate the Lustre and NAS file system infrastructures.

In FY09, LLNL developed and deployed two new Lustre releases that focused on reliability and availability enhancements. Deployment teams also procured, deployed, and operated two major new file systems, including the entire file system infrastructure supporting the Sequoia ID (Dawn). NFS scratch servers were replaced in open computing facility (OCF) and SCF, and home directory servers were replaced in the OCF center.

Planned activities in FY10:

- Maintain GPFS, NAS, and Lustre parallel file system support, middleware, and higher-level I/O library support for users
- Develop and provide Lustre file system performance and scalability enhancements in support of new platforms
- Complete initial Zettabyte File System-based Lustre file system development and testing in preparation for Sequoia petascale environment
- Deploy data analysis NAS file system into production
- Perform parallel file system risk mitigation study

Expected deliverables in FY10:

- Deployment of Lustre release 1.8 into production in OCF and SCF environments
- Deployment of data analysis NAS file system into SCF production
- Report on parallel file system risk mitigation study

Preliminary planned activities in FY11:

- Deploy Zettabyte File System-based file system Lustre release into production in OCF and SCF centers

- Participate in development of the Data Analysis Plan for Sequoia
- Initiate procurement and integration of the file system infrastructure for Sequoia platform

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.

In FY09, LLNL installed the Sequoia ID (Dawn) network. Host adaptor problems were resolved to improve stability of the network.

Planned activities in FY10:

- Perform research and testing for technologies and products for interconnects, local area networks (LANs) and WANs and 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 FY10:

- Switching infrastructure upgrade for future Sequoia network augmentation

Preliminary planned activities in FY11:

- 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
- Test 10GE packet pacing to mitigate packet loss in both LAN because of high CPU usage and WAN because of slow down stream cache overruns
- Continue network improvements and track emerging network technology
- Prepare for procurement and initial installation of Sequoia network

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

The File Systems and I/O Project provides end-to-end, high-performance networking and scalable I/O infrastructure for the ASC program. It also delivers high bandwidth, low-latency interconnect technologies for the ASC compute platforms. The ASC program requires system and storage area network bandwidths at over 500 GB/s., global file system I/O rates beyond 500 GB/s., 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 still a critical concern for current and next-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.

This project is a coordination point for planning of all online storage, network, and data movement activities within the ASC program at LANL. These capabilities include online file systems such as the NFS complex and 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.

In FY09, LANL added several new capabilities: deployed I/O and storage infrastructure for Roadrunner Phase 3; assisted users with I/O issues on Roadrunner phase 3 for open science runs; worked with SNL counterparts to formulate requirements and help evaluate file system and I/O candidate solutions for Mesa; made significant progress on enhancing performance of small and unaligned I/O by writing and testing of the PLFS; and provided initial demonstration of a single name space NFS service.

Planned activities in FY10:

- Support deployment of Roadrunner in the secure by testing the data path to storage including the I/O nodes, PaScalBB infrastructure, and storage
- Continue testing and deployment of LANL's PLSF System on Roadrunner
- Support deployment of Mesa at LANL; in partnership with SNL, work with application developers to optimize I/O performance on this new machine
- Test and stress a single name space NFS service
- Expand the PaScalBB network architecture in the secure to 12 lanes in preparation for Mesa

Expected deliverables in FY10:

- I/O infrastructure preparation and testing for Mesa
- Deployment of PLFS on Roadrunner
- Expanded network capacities in the secure for Roadrunner and Mesa

Preliminary planned activities in FY11:

- Support deployment of enhanced scalable metadata operation
- Evaluate quality of service performance guarantees for file systems

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, 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 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 PFTP, client application programming interface, local file mover and third party SAN (SAN3P). HPSS also provides hierarchical storage management services for IBM GPFS.

In FY09, LANL added the following new capabilities: PSI 2009 release, requirement documents for HPS 7.2 and HPSS 8.0 release 7.2 of HPSS, and the FY10 release of HPSS 8.0.

Planned activities in FY10:

- Finalize requirements for 2010 release of PSI
- Design, develop, and test 2010 release of PSI
- Develop requirements for 2011 release of PSI
- Finalize requirements for release 8.0 of HPSS
- 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 FY10:

- PSI 2010 and 2011 requirements documents
- PSI 2010 release
- Design documents for HPSS 8.0

Preliminary planned activities in FY11:

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

WBS 1.5.4.5-SNL-003 Archival Storage

The Archival Storage project represents SNL's participation in the DOE HPSS Consortium development project. 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.

In FY09, SNL completed and installed parallel file transfer protocols with a HPSS transfer agent. This software increased transfer rates to archival storage from 30MB/s to 200MB/s. SNL also completed and released HPSS version 7.1, which added support for Kerberos authentication, OpenLDAP (Open Lightweight Directory Access Protocol) directory services, and parallel data transfer.

Planned activities in FY10:

- Develop and test security, data-transport performance, and reliability improvements for current and prospective version of HPSS
- Support parallel file transfer protocol/transfer agent/LDAP/security for HPSS

Expected deliverables in FY10:

- Coding support for Version 8.1 of HPSS
- Requirements definition for next HPSS upgrade

Preliminary planned activities in FY11:

- Support development and deployment of current HPSS

WBS 1.5.4.5-SNL-004 Scalable Input/Output Research

The Scalable I/O Research 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 non-volatile storage devices and data-warehouse appliances to improve I/O; and research ways to use available system 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 FY10 R&D 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 and service guarantees, will create new opportunities to push state-of-the-art in I/O.

In FY09, SNL completed design and implementation of GP-graphical processing unit (GPU) enhancements for RAID-6 on commodity hardware and demonstrated 1 GB/s. performance for writes (a five-fold improvement). This advance specifically benefits write-intensive workloads that are common in HPC systems. The software is currently being integrated into the Linux software RAID layer for general distribution. SNL also completed initial implementation of parallel netCDF caching service and demonstrated 150 GB/s. effective write performance (a 10-fold improvement) for netCDF-based I/O benchmark codes. The caching service uses memory on available compute nodes to stage "bursts" of data for write-intensive applications. The netCDF caching service was specifically designed to reduce I/O overhead for application-directed checkpoints.

Planned activities in FY10:

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

Expected deliverables in FY10:

- Parallel file system integrated with cooperative caching in the compute-nodes; the initial implementation will leverage current work using Plan9 on the IBM BlueGene/P system
- Fragment detection service deployed as an in-situ processing service ASC platforms
- GPU enhancements to support on-the-fly compression to back-end storage

Preliminary planned activities in FY11:

- Develop and deploy advanced architecture enhancements into production I/O systems
- Address I/O scalability and reliability issues on Mesa

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 FY10

- Data analysis and visualization software suite deployed on LLNL cluster
- Version 2.4 and 2.5 of LLNL Hopper and Chopper tools
- Visualization-on-platform milestone completion
- A tighter integration of visualization, analysis, and ray-tracing based rendering for improved performance, memory usage, and functionality

- Optimization of end-to-end performance, including delivery of visualized results to user desktops
- Development of readers for EnSight and ParaView for applications of interest to ASC program
- New weapons science visualizations for the Thermonuclear Burn Initiative (TBI) program
- Continue partnership with V&V program to provide advanced tools for visual analysis of complex data, increase partnership with SIERRA codes to provide coupled (in-situ) scalable analysis and visualization capabilities, targeted at capability runs on Mesa, and ongoing technical and user support, including direct analysis support for capability class systems

WBS 1.5.4.6-LLNL-001 Scientific Visualization

The Scientific Visualization project conducts research and develops and supports tools for managing, visualizing, analyzing, and presenting scientific data. Research includes topological analysis, particle visualization, and data compression techniques. Operational support for data analysis covers support of post-processing resources, including visualization servers, displays, software, and facilities. The visualization hardware architecture team engages in planning, test bed prototyping, testing of systems and components, and procurement and integration of new systems. The operational team's server efforts include system administration, computer security, troubleshooting, and maintenance of hardware and software. Display efforts include support of high-resolution, high-performance display devices for theaters and collaborative use areas. Operational support manages theater and PowerWall facilities and associated servers, runs video production labs, and consults on software such as 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. A video display infrastructure drives power walls and smaller displays. The project installs, maintains, and consults on software visualization tools, and supports demonstrations on the PowerWall. The project maintains unclassified and classified video production labs with video editing, 3D modeling and animation tools, and video peripherals for DVDs and videotapes.

In FY09, LLNL supported ASC scientists through visualization and video support and released upgraded versions of movie player software as well as enhancements to other tools in the visualization software suite. LLNL upgraded PowerWall capabilities in two of the visualization theaters, upgrading modems for one and replacing projectors for another.

Planned activities in FY10:

- Perform R&D in topological analysis, data compression, and particle visualization, and mentor students in these areas
- Support petascale data analysis hardware upgrade and integrate a new PowerWall driver and test bed platform
- Support petascale data analysis activities through software consulting

- Maintain and enhance the existing suite of tools and libraries developed by this project in support of PowerWall and visualization cluster usage
- Provide operational support for all visualization facilities, including supporting projection equipment, performing color alignment and projector alignment, and facilitating the use of the data analysis clusters and associated storage
- Support ASC scientists through visualization and video efforts, including the support of PowerWall presentations, creation of visuals and movies to support the presentation of scientific data, and general user support

Expected deliverables in FY10:

- Enhanced remote rendering options
- Upgraded versions of PowerWall movie player software
- Suite of analysis and visualization software deployed on new SCF interactive data analysis cluster

Preliminary planned activities in FY11:

- Continue to enhance and maintain the data analysis environment in both hardware and software areas
- Leverage research efforts in data management and analysis and perform new research in topological analysis for scientific computing and data compression

WBS 1.5.4.6-LLNL-002 Scientific Data Management

In the age of petascale computing environments, the complexity and scale of the data management challenge is also reaching new heights. The Scientific Data Management project provides users with powerful and time-conserving 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.

In FY09, LLNL released versions 2.2 and 2.3 of Hopper and Chopper, featuring built-in file comparison capabilities, support for Livermore Computing's (LC's) centralized storage cluster architecture, an enhanced directory viewer for quickly surveying the contents of a directory, support for direct-to-tape heuristics, and improved scalability.

Planned activities in FY10:

- Maintain and enhance the existing suite of tools and libraries developed by the scientific data management project

- Investigate additional levels of concurrency, for example, for searches and recursive directory operations

Expected deliverables in FY10:

- Version 2.4 of Hopper and Chopper, featuring a graphical disk usage view and associated framework, enhanced graphical previews for various file types, a new interface to the file exchange mechanism, and optimization of long-running directory operations.
- Version 2.5 of Hopper and Chopper, featuring full desktop support of indexed tar files, optimization of multi-TB tar/htar operations, toolbar editor, and support for user-provided extensions to the launch facility

Preliminary planned activities in FY11:

- Continue to maintain and enhance the Scientific Data Management suite of tools and libraries based on user feedback and requirements imposed by a petascale computing environment
- Explore means for retrieving disk usage information more efficiently from massive file systems, such as Lustre and HPSS

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

The Visualization and Insight for Petascale Simulations Project 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.

The project focuses on 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 one can analyze as the data is computed, one can avoid writing to disk and then re-reading the data); and 2) for performance, using the power of the petascale platform to analyze petascale data.

In FY09, LANL enhanced integrated distance and comparative/feature extraction visualization capabilities in software tools, wrote an implementation plan and implemented accelerated visualization using cell or other technologies in user tool base, and developed readers for Ensight and ParaView for the VPIC code.

Planned activities in FY10:

- Develop visualization, analysis, rendering software aligned with emerging supercomputing platforms to advance petascale data analysis infrastructure
- Explore and evaluate the maturity level of visualization and rendering on the supercomputing platform including scheduling and desktop delivery of results

Expected deliverables in FY10:

- Completion of the visualization on the platform milestone, to support the JEI, by delivering evaluation results documented as a report and tools made available to the ASC community
- Research on new petascale visual analysis approaches is documented, including statistical, feature, and analysis operators
- Readers for Ensight and ParaView developed for applications of interest to ASC program

Planned activities in FY11:

- Transfer technology to the production environment, based on a positive evaluation of the FY10 Level 2 milestone *Visualization-on-platform technology*
- Explore the end-to-end integration of in-situ analysis, feature-extraction, and ray-tracing based visualization to improve the understanding of massive petascale results

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 through six sigma processes.

Visualization and visual analysis are essential tools needed by code teams and designers in understanding the terabytes of data that are generated in a single simulation run. This project, also referred to as "Production Visualization," provides visualization services from the machine to the desktop for users in the ASC program. People funded by Production Visualization also 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 ASC 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.

In preparation for exascale, it will be necessary to analyze coming computer systems and how they may be appropriately used for visualization and data analysis applications. It will also be necessary to examine infrastructure and redesign and deploy appropriate infrastructure as needed. Finally, display solutions, designed in the terascale era, will need to be revisited and redesigned as needed. New interaction paradigms may also be appropriate.

Capabilities include custom ASC visualizations; support of the large facilities, such the CAVE and the Powerwall Theater, including assistance, training, and developing new tools to work with these facilities; briefing support; and deployment and maintenance of any needed Visualization Corridor software.

In FY09, LANL kept facilities up nearly 100 percent of the time and used them for user visualization, director-level briefings (such as for Secretary Chu), weapons program

seminars, reviews, colloquiums, and programmatic collaborations. LANL started early planning for Viewmaster upgrades and stood up an open ASC visualization resource from the older Viewmaster R&D nodes. This is the first open ASC visualization resource at LANL in five years.

Planned activities in FY10:

- Continue the support and maintenance of the large visualization facilities
- Conduct briefings as needed
- Support users in providing visualizations of simulations, experimental data, and engineering computer-aided design (CAD) data on desktops and in facilities
- Plan and execute Viewmaster follow-on to support hardware-enabled visualization
- Help direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International

Expected deliverables in FY10:

- High availability in the large visualization facilities
- Scientific visualization support for briefings
- User visualizations

Preliminary planned activities in FY11:

- Continue to support and maintain the large visualization facilities
- Continue to support users in providing visualizations of simulations, experimental data, and engineering CAD data on desktops and in facilities

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, thus improving 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 analysis problems and to promote new weapons science discoveries using the ASC codes.

This small group of individuals has played a major role in the recent successes of the TBI program at LANL. Their work in combining ASC visualization with the new capabilities provided by the ASC codes has been one of the major factors responsible for the new discoveries being made by the TBI program.

In addition to working directly with the design community on its visualization and analysis problems, this group is also responsible for some 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, Inc. for problem reporting and resolution and for new feature requests. This group, with the Production Visualization group, also directs all subcontracts that LANL has with Computational Engineering

International, Inc. related to new EnSight development and to onsite training and consulting.

In FY09, LANL added the following new capabilities: visualizations of a NIF inertial confinement fusion simulation; new weapons science visualizations for the TBI program; other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs; and PoP and General Mesh Viewer with modest updates to fix bugs and improve portability and maintainability.

Planned activities in FY10:

- Work directly with designers in physics-based, iterative discovery process using a petascale visualization and data analysis enabled tool (EnSight)
- Support and maintain the EnSight software and help direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International
- Document work with joint publications co-authored with X Division designers on a variety of weapons science topics

Expected deliverables in FY10:

- New weapons science visualizations for the TBI program
- Other visualizations and data analysis products to address a variety of DSW-related activities, including the closing of outstanding SFIs

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 TBI
- Continue to document the results of 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 focuses on providing advanced customer-centered capabilities within an open source production framework, enabling those capabilities to be utilized by a variety of end user tools, specialized for many problem domains. Foundational capabilities are released in the open source Visualization Toolkit, which serves as the foundation for a number of scalable tools. LANL 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. The 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.

This project also provides deployment and support services that enable ASC customers to carry out data analysis on Mesa and future systems. This includes providing technical assistance for capability runs on the ACES platforms, as well as bridging the gap between advanced R&D and the users who must work with the entire end-to-end modeling and simulation environment.

In FY09, SNL added two new capabilities: 1) crater analysis capability that leverages the scalable scientific and information visualization capabilities within ParaView to enable analysts to have a broader understanding of ensembles of runs by providing additional

quantitative data from a simulation run, including abstract features like fragments and craters; and 2) initial implementation of in-situ analysis capability, which has been tested in conjunction with several codes (as called for in the *ASC Infrastructure for Petascale Environments* document). This capability provides a general analysis toolset (including the fragment and crater characterization work from milestones in FY08 and FY09) that can be used in conjunction with a running simulation to provide a rich set of data output, including statistics, simplified geometry, sparse variable output (including newly calculated variables), images, and movies. This capability will be especially useful as the move to Mesa and beyond continues.

Planned activities in FY10:

- Continue partnership with SNL's V&V program to provide advanced tools for visual analysis of complex data
- Partner with SIERRA codes to provide coupled (in-situ) scalable analysis and visualization capabilities, targeted at capability runs on Mesa
- Provide ongoing technical and user support, including direct analysis support for capability-class systems

Expected deliverables in FY10:

- Level 2 milestone to provide advanced visualization capability in support of FY12 V&V milestone
- Joint Level 2 milestone with LANL to provide initial implementation of advanced rendering capabilities in support of Mesa
- Multi-threaded and multi-core re-architecting of Visualization Toolkit and ParaView, to support ACES platforms, including Mesa
- Advanced analysis and visualization library for direct coupling with codes
- Ongoing releases and installations of ParaView

Preliminary planned activities in FY11:

- Deliver advanced analysis capability in support of GA of Mesa, 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

WBS 1.5.4.7: Common Computing Environment

The goal of the 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 FY10

- Tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"
- TOSS 1.3 (Based on RHEL 5.4)
- O|SS design study for distributed and componentized tool infrastructure and test results for selected individual tool components, such as data collection and instrumentation, outside the current framework
- Expanded WC Tool that includes usage reporting with a reporting interface that consolidates all usage and machine statistic reports
- Continued tool and user support for the production version of application monitoring tool MoJo running on TLCC platforms in the tri-lab community
- Application monitoring and analysis tool prototype with a common, scalable, data organization and access scheme
- Shared work space tri-lab GForge capability for project planning and repositories
- Demonstration of Gazebo Test and Analysis Suite Web-based interface working with other monitoring applications
- Demonstration of STAT debugger functionality on tri-lab machines and provide a common approach for scalable subset debugging
- Requirements, plans, results for security integration based on scope

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 CCE requirements for providing 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 a production-quality software environment can be deployed across the tri-lab in a consistent and manageable fashion.

In FY09, TOSS 1.2 (based on RHEL 5.3) production version was released for GA on TLCC systems. Additional FY09 deliverables included 1) 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; 2) reasonable support processes including active and consistent issue tracking; and 3) appropriate coordination amongst the labs for overall TOSS support (i.e. support meetings happening on regular intervals).

Planned activities in FY10:

- Provide ongoing TOSS software development and support
- Develop/Deploy TOSS 1.3 (Based on RHEL 5.4)

- Develop beta release of TOSS 2.0, contingent on RedHat RHEL 6 schedule
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Deployed TOSS 1.3 (Based on RHEL 5.4)
- Delivered beta release of TOSS 2.0, contingent on RedHat RHEL 6 schedule

Preliminary planned activities in FY11:

- Provide ongoing TOSS software development and support
- Deploy GA release of TOSS 2.0, contingent on RedHat RHEL 6 schedule
- Prepare for deployment of the next generation of the ASC TLCC systems, which may include software integration and testing for the tri-lab environment

WBS 1.5.4.7-TRI-002 Open|SpeedShop

O|SS is targeted to be the main performance analysis tool set across all tri-lab ASC production systems. 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. It is currently available on most TLCC installations as well as selected other capacity cluster architectures.

While the development of the O|SS infrastructure for capacity clusters is nearing its end and will in the near future focus on new tool functionality as well as maintenance and support for the overall tool, new development efforts, funded by this project, will focus on extending O|SS to capability machines, including the Sequoia ID (Dawn), Sequoia, and Mesa.

To address these challenges, NNSA has partnered with the Office of Science/Office of Advanced Scientific Computing Research (ASCR) on extending O|SS to NNSA's flagship machines as well as to ASCR's LCF architectures. This funding represents NNSA's contribution to the overall project and ensures that extensions and redesigns of O|SS match the requirements for ASC's machines. Additionally, this project includes personnel at all three labs to provide technical guidance and ensure early adoption of O|SS across ASC platforms.

In FY09, O|SS was deployed at all three labs. New capabilities delivered in FY09 include latest software release 1.92; enhanced offline performance monitoring leading to increased stability and production use; increased scaling capability (2048 processor runs can now be done routinely); and new short cut startup commands as requested by tri-lab user community. Additional FY09 lab-specific accomplishments include:

- LLNL: continued support and user interactions on TLCC machines, O|SS is available through dotkit infrastructure. Documentation on the LC Web pages has been added. Additionally, O|SS is installed on the CASC research cluster for easier interaction with other ASC and LLNL researchers. This will help extend the capabilities of O|SS. Extended testing and successful runs of O|SS were completed on the LLNL Hyperion test bed.

- LANL: established O|SS as a module on all machines. A quick-start guide is available through LANL HPC support Web pages.

- SNL: established as a module on TLCC machines. Verification testing done on research machines. Additionally, O|SS is installed on the Catalyst cluster for testing and research.

Planned activities in FY10:

- Execute application studies with ASC and ASCR code teams to gather requirements and design guidelines for petascale performance analysis
- Redesign the O|SS infrastructure into an open, distributed, and componentized tools platform
- Evaluate design studies on current capability platforms
- Integrate existing internal and external tool components (including Stackwalker and Launchmon) to improve code reuse, tool portability, and stability
- Support the tri-lab CCE Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Design study for distributed and componentized tool infrastructure
- Test of selected individual tool components, such as data collection and instrumentation, outside the current framework
- Initial prototype port to BlueGene architecture
- Continued production support on capacity machines

Preliminary planned activities in FY11:

- Continue the three-year project in collaboration with ASCR (FY10–FY12)
- Introduce distributed data preprocessing, aggregation, and analysis
- Deliver prototypes of a highly distributed and concurrent version of O|SS capable of scaling to petascale machine sizes

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) includes: developing new functionality in a WC Tool; modifying existing local laboratory tools; and integrating the WC Tool with the Moab resource manager and local laboratory tools and databases.

The WC Tool was developed based on lessons learned from SNL's Web-based HPC Estimations and Requirements Tool and more than two years of experience with the tool. The WC Tool provides the capability to collect and report current and future requirements for compute resources, with programmatic characterization of the work; and computing resource usage, with programmatic characterization of the work. The tool includes a modular mechanism that can interface to multiple existing databases at each of the tri-labs. Additional development is needed in the areas of validation, test suites, and common reporting capability.

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

The Workload Characterization project will expand in FY10 to include additional computing resource usage and machine utilization reporting. This work will include integration with the underlying job databases (for example, SLURM accounting database).

In FY09, the project developed a tri-lab common reporting interface, the WC Tool, to collect, analyze and report on compute resource requirements for current and future use, with programmatic characterization of the work; and platform usage data, tied to the programmatic characterization of the work. For the first time, the tri-labs are able to report ASC computing resource demand and usage with programmatic characterization of the work.

Planned activities in FY10:

- Enhance the WC Tool to report additional job usage and machine utilization
- Gather requirements from NNSA/HQ and each laboratory to expand the reporting interface to generate reports of additional computing resource usage and machine utilization
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Expanded WC Tool with additional computing resource usage reporting
- Reporting interface that consolidates computing resource usage and machine statistic reports

Preliminary planned activities in FY11:

- Assess usage and provide enhancements in tools or documentation and training
- Improve the WC Tool, in coordination with the other laboratories in the context of the CCE projects

WBS 1.5.4.7-TRI-004 Application Monitoring

The Application Monitoring project is developing a basic set of monitoring tools, along with their system and application interfaces to facilitate automated monitoring of production ASC systems and applications. The tools provide information about user applications, system, and component state to answer questions about progress, interruption, intervention, and notification. The result will be an extensible framework for future application and system monitoring features.

The MoJo tool is the project's prototype implementation developed in FY08 and FY09 to address the application part of this. MoJo uses a centralized database to monitor application progress. The initial implementation of MoJo can identify progress through file system interactions (files being created or modified) at a user-specified interval. This is being expanded in FY10 to address tying this together with hardware and system monitoring for an end-to-end solution.

New capabilities added in FY09 include development of MoJo, based on combination of related works from SNL (JobMonitor) and LANL (job_suite) with input from LLNL;

deployment of MoJo on TLCC machines at SNL and initial deployment at LANL; demonstration using the MPI Profiling layer to gather statistical data to prototype advanced MoJo features in a tool called Easy Monitoring, a lightweight, generic monitor/logger for MPI programs; and prototype interface between MoJo and a system monitoring tool at LANL.

Planned activities in FY10:

- Develop a common, scalable, data organization and access scheme that will facilitate flow of information between application, system, and hardware platform critical to increased application efficiency.
- Determine and prototype desired MoJo enhancements and advanced features including integration of above scheme
- Design standalone version of MoJo that does not require a centralized database

Expected deliverables in FY10:

- Prototype system employing a common scalable, data organization and access scheme, and APIs to share appropriate data between a sample application, resource manager, and hardware platform
- Continued tool and user support for the production version of MoJo running on TLCC platforms in the tri-lab community
- Prototype implementation of MoJo with advanced functionality that includes additional application data (list assembled at FY09 workshops)
- Standalone version of MoJo that allows for data gathering in environments where a centralized database is not possible

Preliminary planned activities in FY11:

- Deploy prototype of system using capabilities developed in FY10 on a production TLCC system to assess productivity gains
- Deploy production version of advanced MoJo functionality

WBS 1.5.4.7-TRI-005 Shared Work Space

The Shared Work Space 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.

In FY09, Shared Work Space was deployed in a limited capability. It includes 1) Gforge server with 100 user licenses (this is the software that provides the ability to create project spaces and includes tools for project management, document management, code repository and other features http://gforge.org/gf/); 2) tri-lab deployment through a server running on SNL's restricted network; 3) tri-lab access through cross-realm authentication, which allows a user at labs external to SNL to utilize their respective authentication means; and 4) limited capability due to security access issues that were identified in regard to email communication and repository access in the current

network and security plan rules. It has been deployed to two cross-realm teams to use in the interim. A new server on SNL's extranet should resolve security access issues.

Planned activities in FY10:

- Support the GForge server, licenses, and continued tri-lab planning for enhancing communication as the Shared Work Space project moves into production in FY09 (the GForge server is housed and managed at SNL; GForge admin support for the environment is done through an identified person at each lab; licenses are purchased through Los Alamos)
- Define needs through interaction with the Security Integration project
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Tri-lab GForge capability for project planning and repositories, Wikis
- Additional review of tools that support tri-lab communication

Preliminary planned activities in FY11:

None (this project is completed in FY10)

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 that interrogate and report the status of individual hardware components comprising the cluster.

In FY09, Gazebo was updated and packaged as an SRPM; Gazebo SRPM was rebuilt into RPM for inclusion in TOSS stack. In early FY09, Gazebo was used extensively by the tri-labs as a component of the SWL acceptance testing for all the TLCC systems. Also in FY09, Gazebo work focused on development of a graphical analysis tool and the integration of the CBENCH tool suite into Gazebo.

Planned activities in FY10:

- Provide enhancements to Web-based reporting tool
- Integrate Web application with other monitoring project applications
- Participate with design and development of common monitoring project schema definition

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

Expected deliverables in FY10:

- Gazebo Web-based interface working with other monitoring applications

Preliminary planned activities in FY11:

- Integrate with all CCE projects under common authentication and authorization environment

WBS 1.5.4.7-TRI-007 Debugger

All three labs are using TotalView as the core debugger. LLNL has a more strategic relationship with TotalView Technologies and is also performing additional research into other debugging tools (for example, STAT). The approach to scalable subset debugging is key as ASC moves toward larger scales and collects information that helps focus in on the trouble area. Input from all three labs and target applications are required. The working group would like to build a tri-lab capability around LLNL.

As a newly started project for FY10, a new sub-team will be created for communication and support. A capability will be developed to provide cross-lab trials for STAT with a focus on subset debugging that would also utilize TotalView and support scaling efforts. The University of New Mexico will provide technical guidance on porting and adapting the base infrastructure. A central integration point for TotalView strategic planning will be created through LLNL.

Planned activities in FY10:

- Develop a more cohesive debug capability aligned with the CCE goals of commonality and leverage of infrastructure and plans
- Address debugging on capability machines through the scalable-subset-debugging segment of this effort
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Joint tri-lab planning sessions with TotalView
- STAT on tri-lab machines and a common approach for performing scalable subset debugging

Preliminary planned activities in FY11:

- Tri-lab deployment of enhanced functionality developed in STAT

WBS 1.5.4.7-TRI-008 Security Integration

With the advent of the TLCC effort, LANL, LLNL, and SNL have an increased need for integrated access to unclassified resources for improved collaboration and to better utilize computing resources. Attempts to facilitate such collaborations have been hampered by incompatible security policies, interpretations, and implementations that make it difficult to access unclassified resources cross-site. This activity is a component of the multi-site security integration effort.

A key capability is to identify and implement a tri-lab security approach to achieve the following targets:

- Security policies and implementations that allow tri-lab access using identified protocols and technologies
- More effective resource control and utilization
- Utilization, where feasible, of home site security apparatuses, including authentication equipment (for example, single sign on)
- User access to tri-lab resources within a specified period of time upon request based on identified need
- Collaboration with the CCE user community to develop security policies that facilitate easier resource utilization
- Collaboration with the appropriate security personnel to ensure timely implementation
- Support of multiple classes of users such as architecture and modeling, application, customer service, system administrators with root access, and end users
- Necessary network bandwidth and latency to cover current and projected requirements

Planned activities in FY10:

- Support security integration architecture team in defining needs and developing approaches to solve tri-lab security issues
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Requirements, plans, and deliverables based on identified scope

Preliminary planned activities in FY11:

None

WBS 1.5.4.7-TRI-009 Open Source Contract Maintenance

This past year ASC assessed the model for open source tool support and working with open source communities for better sustainability. An outcome was to create support contracts for these communities/organizations through the CCE for long-term tri-lab support. The benefits—as opposed to individual lab contracts—are that the tri-lab coordinates its needs better, has a single focused contract, and can possibly take advantage of the contract model at LLNL for cost savings.

A process/capability was established for support contracts for open source products at a tri-lab level. Note that these funds are not intended for development of new functionality but rather for long-term maintenance work and support activities. The tri-lab will continue to work with other open source organizations to establish models that are beneficial to the CCE.

Planned activities in FY10:

- Support O|SS, Krell Institute
  - Code team interactions (support calls/work)

  - Training and tutorials at all three laboratories
  - Testing on lab platforms, release validation on target lab systems
- Support Valgrind, OpenWorks
  - Support for tri-lab platforms (capacity/capability)
  - Release validation against MPI and compilers on TLCC and other systems
  - Expedited bug resolution and patch releases for the tri-labs
- Support other efforts as required
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- Effective support structure for target products

Preliminary planned activities in FY11:

- Continue this effort towards a long-term support model

WBS 1.5.4.7-TRI-010 OpenMPI Integration/Scaling

OpenMPI is used as the primary MPI at LANL and SNL and as the secondary MPI at LLNL. There is a continued need for better integration into the TOSS stack and scaling support tied to the resource manager product. LANL is a member of the OpenMPI development community and can effect capability direction. The addition of tri-lab testing and support will allow better application efficiency through parameter studies on TLCC clusters against various application usage models.

A capability was established to support better integration to the SLURM segment of the TOSS stack and to provide tri-lab support for issues and parameter analysis for environment differences. The project also supports the CCE goal of better integration of runtime tools; the scaling studies and work targeted toward tighter integration to the resource manager will provide faster MPI startup times targeted to capability machines

Planned activities in FY10:

- Perform scaling work with resource manager (LANL) for both capacity and capability computing
- Initiate cross-lab parameter studies, testing and support (LANL)
- Support the tri-lab Level 2 milestone "Deploy CCE capabilities for capacity computing environment"

Expected deliverables in FY10:

- OpenMPI enhancements for better scaling
- Process for assessing parameter identification

Preliminary planned activities in FY11:

None

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 FY10

- Interactive data analysis cluster system integration and support
- Integration of Sierra capacity cluster
- NNSA Policy Letters (NAPS)-compliant security plan for the LC
- One-Way Link File Interchange System deployment to enable file transfers from unclassified to classified network
- Activation testing of the electrical components from the first phase of construction of the B-453 15-megawatt power expansion (7.5 MW west room electrical equipment)
- Document containing the results gathered during networking technology evaluations and a recommended architecture for Sequoia deployment
- SCC infrastructure upgrade project, which will bring the power to the computer room floor to 16.8 MWs
- Site preparation and infrastructure upgrades for delivery of Mesa
- Replacement of Laboratory Data Communications Center (LDCC) cooling tower pumps
- Power/cooling metering installation
- Analysis and solutions to reduce power usage effectiveness in LDCC and SCC computing facilities

- Involvement in the procurement process for Mesa and maintenance contracts for hardware/software for this platform
- Deployment of the FY09 tri-lab common software environment on the production TLCC systems
- Delivery of classified and unclassified capability and capacity computing cycles for the tri-lab
- Development and delivery of enhancements to a common report format for metrics on ASC platforms at LLNL, LANL, and SNL

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; and 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.

In FY09, LLNL moved 40 racks of BlueGene/L to the open side, creating an unclassified resource of greater than 200 teraFLOPS for unclassified science. The integration of the Sequoia ID (Dawn) system was completed. Lilac and ALC Linux clusters were retired and replaced with 18 TLCC SU: Juno, Eos, Minos on the classified side, and Hera on unclassified side.

Planned activities in FY10:

- Move the Sequoia ID (Dawn) to classified network and attain GA
- Integrate interactive data analysis cluster and achieve production status on classified network
- Integrate Sierra, a capacity cluster jointly owned by LLNL and ASC
- Support the Sequoia ID (Dawn), Purple, BlueGene/L, Peloton, and TLCC capacity systems

Expected deliverables in FY10:

- Interactive data analysis cluster system integration and support
- Sierra capacity cluster integration and support
- Ongoing support of Purple, the Sequoia ID (Dawn), BlueGene/L, Peloton, and TLCC capacity systems
- Ongoing support of Purple
- Preparations for deployment of Sequoia

Preliminary planned activities in FY11:

- Support the Sequoia ID (Dawn), Purple, BlueGene/L, Peloton, and TLCC capacity systems
- Deploy 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.

In FY09, LLNL completed the negotiations related to the Sequoia ID (Dawn) system and began self-maintenance. This included spares inventory and operations staff training and documentation. LLNL also tracked and placed contracts and licenses needed for system operations and vendor support; performed ongoing hardware self-maintenance of BlueGene/L, parallel global file systems, Peloton, and TLCC capacity systems.

Planned activities in FY10:

- Track and place contracts and licenses needed for system operations and vendor support
- Provide ongoing hardware self-maintenance of the Sequoia ID (Dawn), BlueGene/L, visualization clusters, parallel global file systems, Peloton, and TLCC capacity systems

Expected deliverables in FY10:

- Contracts and licenses needed for system operations and vendor support
- The Sequoia ID (Dawn), BlueGene/L, visualization clusters, parallel global file systems, Peloton, and TLCC capacity systems properly maintained
- Self-maintenance of the Sequoia ID (Dawn)

Preliminary planned activities in FY11:

- Put in place contracts and licenses needed for system operations and vendor support
- Maintain properly BlueGene/L, the Sequoia ID (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.

This project also ensures that computer security plans meet NNSA-specific requirements.

In FY09, LLNL deployed the commercial identity management solution phase I. Deployed features include account creation, modification, and deletion on both unclassified and classified systems.

Planned activities in FY10:

- Continue upgrades to core security infrastructure components, thus enabling enhanced security capabilities and mechanisms to be offered and utilized
- Integrate security middleware software stack for CHAOS and platform operating system upgrades
- Develop and integrate additional electronic workflow and provisioning functionality that leverages the deployed identity management solution
- Consolidate platforms and systems management for the core security infrastructure in the unclassified and classified computing environment
- Participate in the effort to build a Site Security Component Library, including preparing security test and evaluation plans
- Perform ongoing security-related activities in support of the secure use and secure management of ASC platforms and associated infrastructure
- Develop the NAPS-compliant security plan for the LC

Expected deliverables in FY10:

- Deployment of identity management phase II features, including privileged account management, bulk operations and electronic provisioning for the classified environment
- Approval to test under the delivered NAPS-compliant LC Information System Security Plan
- Upgrade of RSA authentication manager to the latest stable release
- Development of capability to support data encryption for off-site disaster recovery backups
- Deployment of a one-way link file interchange system to enable file transfers from unclassified to classified network
- Transition of the user base on the classified computing environment to use two-factor authentication for all logins
- Porting of authentication and security registry services to utilize a common platform

Preliminary planned activities in FY11:

- Integrate the data management for RSA authentication manager with security registry and identity management services.
- Continue to incorporate external data sources and evaluate process improvements for a comprehensive identity management system

- Integrate security middleware software stack on next generation TLCC systems

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.

In FY09, LLNL completed the west room 7.5-MW electrical distribution expansion from the electrical yard into the first level of the machine room.

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-MW electrical power expansion for B-453 for 2011 completion
- Commission and complete activation testing of the electrical components from the first phase of construction of the B-453 15-MW power expansion (7.5MW west room electrical equipment)
- Start the second phase of construction of the B-453 15-MW power expansion (7.5 MW east room electrical equipment)
- 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 implement the results from self-benchmarking tool for all of the computer rooms in B-453 (the Terascale Simulation Facility, TSF), B-451, B-439, B-115, and B-117 created by DOE Office of Science and Lawrence Berkeley National Laboratory (http://hightech.lbl.gov/datacenters.html) to continue to routinely identify prospective energy savings initiatives as computer rooms change
- Evaluate the implementation of free cooling to improve power usage effectiveness from 1.34 to 1.2
- Continue to update computational fluid dynamics model for all ASC systems to profile the airflow required to cool the machines adequately as platforms are added and retired

Expected deliverables in FY10:

- Site preparation studies for Sequoia infrastructure planning
- Construction of the second phase of the B-453 15-MW expansion
- Final cooling tower cell installation

Preliminary planned activities in FY11:

- Begin construction of site preparation for Sequoia infrastructure
- Install final four air handlers in the mechanical rooms
- Continue to maintain and support the equipment in existing computational and staff facilities

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.

In FY09, LLNL removed the Building 451 classified network, upgraded the Weapons & Complex Integration Directorate closed Labnet connections to 10 gigabit, and deployed a firewall in place of router access control. An institutional classified boot network was designed and deployed.

Planned activities in FY10:

- Evaluate the use of Quad Data Rate InfiniBand as a storage area network for the global file system
- Evaluate 40 Gb/s. data center Ethernet as a storage area network for the global file system
- Evaluate use of 40 gigabit Ethernet connectivity in the backbone
- Provide network architecture design and testing to support the Sequoia platform
- Provide ongoing local network performance, reliability, security, and operational support

Expected deliverables in FY10:

- Document containing the results gathered during the networking technology evaluations and a recommended architecture for Sequoia deployment

Preliminary planned activities in FY11:

- Provide ongoing local network performance, reliability, security, and operational support
- Continue to analyze emerging network technology

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.

In FY09, LLNL began deployment of 10-gigabit encryption for the DisCom WAN.

Planned activities in FY10:

- 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 FY10:

- Ongoing operational support of the tri-lab WAN connections

Preliminary planned activities in FY11:

- 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.

In FY09, LANL deployed a common workload characterization software on platforms running the MOAB resource management tools, and deployed a common report format on platforms running the MOAB resource management tools.

Planned activities in FY10:

- Coordinate with other nuclear complex laboratories to integrate functionality from existing HPC job and machine usage statistics reporting tools into the WC tool. Include as additional functionality the ability to produce reports required by the Office of Management and Budget.
- Develop requirements from each laboratory for an interface that will generate all the reports of computing resource usage, machine, utilization, as well as the actuals versus estimates from the WC Tool.
- Develop and deploy the browser-based user interface as well as command line interface. The reports generated by this tool would be tables, charts, and graphs displayed in the browser, with an option the capability to save as pdf and Excel files. The command line interface would output directly to the command line with an option to output to a file.

Expected deliverables in FY10:

- Reporting interface that consolidates all usage and machine statistic reports

Preliminary planned activities in FY11:

- Respond to laboratory and HQ direction for further refinements to WC Tool, in coordination with the other laboratories in the context of the CCE projects

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.

In FY09, LANL implemented the classified Enterprise SecureNet network, increased network bandwidth to TLCC clusters, and added 20Gbits bandwidth to archival storage and Viewmaster visualization cluster.

Planned activities in FY10:

- 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, reliability, and availability of backbone and services

Expected deliverables in FY10:

- New technology deployment to increase network utilization, bandwidth, and reliability
- Internet protocol encryptor upgrade
- Ongoing network support for all ASC systems
- Implement I/O infrastructure to support Mesa

Preliminary planned activities in FY11:

- Continue to operate and maintain LAN/Metropolitan Area Network/WAN infrastructure
- Continue to integrate with Enterprise SecureNet network

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.

In FY09, LANL successfully deployed network infrastructure for Roadrunner Phase 3; successfully deployed network infrastructure upgrades in support of Hurricane, Turing,

and Lobo TLCC clusters; and developed an initial network plan for Mesa and its file system.

Planned activities in FY10:

- Execute the plan for integrating Mesa into the redesigned I/O network infrastructure
- Execute the plan for integrating Mesa and its infrastructure into to the Integrated Computing Network infrastructure
- Extend the 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, reliability, and availability to HPSS, visualization platforms, and customer workstations
- Provide support for system interconnects monitoring, reliability, and performance management

Expected deliverables in FY10:

- Completed integration of Mesa into the network infrastructure in preparation for full-scale operation
- Network optimization to improve performance to Roadrunner Phase 3 and TLCC clusters
- DisCom encryptor upgrade to 10-Gig Ethernet unit

Preliminary planned activities in FY11:

- Plan I/O infrastructure to support next cluster platform

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.

In FY09, LANL added the following new capabilities: HPSS v7.1 deployment in classified and unclassified networks; production support for Roadrunner Phase 3 system; production support for Hurricane, Turing, and Lobo clusters; monitoring infrastructure deployed on Roadrunner Phase 3 system; and LANL and SNL reviewing best practices in system integration and support via the ACES partnership in preparation for Mesa.

Planned activities in FY10:

- 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
- Improve continuously 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
- 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 data storage systems
- Ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Train operations staff in the hardware maintenance of Roadrunner and TLCC
- Work with system administrators to enhance the monitoring tool infrastructure to include additional features, such as asset tracking and ability to cascade to hierarchical windows for detecting node problems
- Provide hardware support for installation and integration of additional scientific computing platforms, including Mesa
- Extend LANL and SNL best practices in system integration and support via the ACES partnership

Expected deliverables in FY10:

- Support for Roadrunner weapons science and select integrated weapons applications
- Contribution to two new releases of the CCE TOSS software stack
- Monitoring infrastructure deployment to all existing ASC platforms
- Preparation for an initial installation of Mesa
- Hardware support for the installation and integration of Mesa as well as future TLCC systems
- Hardware repair support for Mesa
- Highly improved monitoring system for all ASC platforms
- Monitoring support for all ASC systems, including data storage

Preliminary planned activities in FY11:

- Transition Mesa to full tri-lab capability usage and support
- Support an expanded set of weapons applications on Roadrunner
- Provide hardware support for all ASC platforms

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 FY10 will be the power and cooling infrastructure equipment upgrades to the SCC in support of Mesa. 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.

In FY09, LANL added the following new capabilities: upgraded mechanical and electrical equipment in the SCC in preparation for Mesa, completed final design for the

SCC Infrastructure Upgrade Project for adding an additional 7.2 MWs of raw power and increasing cooling capacity, and replaced 700-ton chillers with three 800-ton water cooled centrifugal chillers in LDCC. Additionally, LANL designed and deployed a chiller water loop.

Planned activities in FY10:

- Upgrade existing mechanical and electrical equipment in the SCC in preparation for Mesa
- Continue operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide facility support for decommissioning of Lightning/Bolt systems
- Increase power usage 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
- Complete SCC Infrastructure Upgrade Project, which will bring the power to the computer room floor to 16.8 MW
- Complete site preparation and infrastructure upgrades for delivery of the Mesa capability system
- Continue involvement in the power and cooling design requirements for the Mesa platform
- Replace LDCC cooling tower pumps
- Assist in the decommissioning and removal of LNXI Lightning/Bolt platforms
- Design chiller water loop in Room 341 (LDCC)
- Install power/cooling metering
- Continue analysis and solutions to reduce power usage effectiveness in LDCC and SCC
- Continue involvement in the procurement process for the Mesa capability system as well as any maintenance contracts for hardware/software for this platform
- Participate in the decommissioning and removal of the LNXI Lightning/Bolt platforms

Expected deliverables in FY10:

- Completion of the infrastructure equipment upgrade in the SCC in support of Mesa
- Site preparation for Mesa in the SCC is complete
- Enhancements to electricity and cooling to support expected increase in computing capacity in the LDCC along with new system cooling requirements
- Retired systems, Lightning/Bolt, are decommissioned

Preliminary planned activities in FY11:

- Begin site preparations for installation of future computing platforms

- Continue operations of the SCC, LDCC, and Central Computing Facility facilities and computer rooms
- Plan infrastructure improvements to keep power and cooling capabilities commensurate with new supercomputer power and cooling requirements

WBS 1.5.5.1-LANL-009 Mesa Capability Computing Platform Initial Deployment

The scope of the Mesa Capability Computing Platform Initial Deployment project is to take delivery and start the deployment of Mesa. This includes completing the acceptance tests, system integration into the LANL network, system stabilization, and transition into the classified network. The primary capabilities are acceptance and diagnostic testing, system stabilization, system integration into the yellow network, and transition to secure network.

Planned activities in FY10:

- Coordinate with selected vendor on the delivery of Mesa at LANL
- Develop acceptance test plan for Mesa
- Deploy initial Mesa and provide system integration
- Create plan for integrating Mesa into secure computing environment

Expected deliverables in FY10:

- Delivery of Mesa system
- Delivery of file system for Mesa
- Evaluation of delivered system with diagnostic testing
- Mesa system integration into LANL network environment
- Completion of system stabilization period
- Completion of security and security test plans

Preliminary planned activities in FY11:

- Complete acceptance testing for Mesa
- Complete transition of Mesa into classified LANL network
- Develop integrated support structure with vendor for operational issues

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 supports tri-lab capability platform resource allocations and coordinates 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 CCE common service and environment decisions and configuration management activities will also be provided.

In FY09, SNL created capability to support DOE/IN and Intelligence Community by converting the previously unclassified head of Red Storm. The new concept, termed Dark Storm, can support several distinct and separate security realms by providing an externalized file system for each specific security enclave. These external file systems are connected individually to Dark Storm as necessary to support simulation runs and analysis needed by the Intelligence Community customers. Customers retain control of their data and share the operating expenses of Dark Storm, using only what they need and timesharing the resource.

Planned activities in FY10:

- Operate production systems in support of nuclear weapons mission requirements—capability, capacity, data services and visualization systems, and long-term hierarchical storage
- Operate three TLCC07 capacity systems: Unity, Whitney, and Glory in production
- Establish with LANL the operational plan for Mesa and document the service level agreement

Expected deliverables in FY10:

- Three TLCC capacity production systems running the CCE software stack
- FY09 tri-lab CCE software environment deployed on the production TLCC systems

Preliminary planned activities in FY11:

- Deploy TLCC11, the next generation of ASC capacity computing systems, if available
- Decommission local dedicated data analysis/visualization platforms (RoSE); replacement as needed with a virtual capability within the capacity computing resources, commensurate with capacity computing needs

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 to conduct secure operations.

In FY09, SNL introduced high efficiency DC powered computer room air conditioners, and implemented a prototype "green energy" effort installing solar panels and a wind power generator to counter higher electricity rates.

Planned activities in FY10:

- Assess and plan acquisitions for incremental power, cooling, and facility modifications to support future-year TLCC installations
- Manage retirement and destruction of classified disk and tape media

- Operate and maintain systems in a fully compliant facility for HPC needs of DOE/NNSA and OGAs supporting the Field Intelligence Elements at SNL and other DOE sites

Expected deliverables in FY10:

None

Preliminary planned activities in FY11:

- 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

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. The project coordinates production requests for tri-lab resources.

In FY09, SNL improved the DisCom WAN reliability by eliminating the one remaining single fiber link of the dual ring network. There are now fully redundant paths connecting the tri-labs, virtually removing any single point of failure locations.

Planned activities in FY10:

- 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 management tools and services between LANL and SNL for Mesa
- Improve monitoring capabilities for data movement tools and services

Expected deliverables in FY10:

- ASC WAN tri-lab capability satisfying FY08 Office of Management and Budget mandate in production

Preliminary planned activities in FY11:

- Provide highly integrated monitoring of local and remote production platforms
- Provide, as needed, and support production tools and remote-usage services for Mesa

WBS 1.5.5.1-KCP-001 Life Extension Program Production Support

The Production Support project directly supports the design/manufacturing issues of current and near-term LEP builds. This effort utilizes the developed software and hardware resources to support production problems and production rate issues at the Kansas City Plant (KCP). This element implements ASC developed codes and capacity HPC hardware for support of final DP weapons production. Current efforts are leveraging prior FY09 work on the SIERRA software framework from SNL and the TLCC capacity platforms from the Nuclear Weapons Complex Tripod effort.

Migration from legacy software tools will continue to the newer framework to support multi-physics application in manufacturing. Continued expansion in areas of fluid interactions, material damage/failure, and uncertainty evaluations in production processes is expected. The work of developing the techniques with the hardware and software tools will be built on the capacity TLCC platforms. Production drivers are the current W76 arming, fusing, and firing production builds and B61 Design and Manufacturability LEP efforts. In conjunction with KCP needs, collaborations will be developed within the Nuclear Weapons Complex production community to exploit needs for the toolset for other applications.

In FY09, KCP expanded the existing TLCC system to another sub-scale platform to handle both "official use only" and secure production application needs. In addition to the hardware, the application software SIERRA has been institutionalized as the lead MPI toolset for production application support in the areas of thermal/structural/dynamic response. This includes welding, encapsulated electronics response, and many other tooling/fixturing requirements of the current LEP systems.

Planned activities in FY10:

- Expand SIERRA toolset to all internal HPC networks for support of shock wave and material failure development
- Collaborate with Savannah River, Y-12, and Pantex on determining physics-based simulation solution needs utilizing SIERRA framework
- Process simulation solutions to key product lines issues at the Production sites using the Sierra/TLCC resources
- Interface with Y-12 on common welding applications and solution techniques

Expected deliverables in FY10:

- Support for production issues on the W76 arming, fusing, and firing non-nuclear LEP
- Development of a methodology for utilization of hardware resources to support B61 efforts
- Determination of physics-based simulation needs at all production sites

Preliminary planned activities in FY11:

- Expand sub-SU TLCC at KCP for B61 development production support
- Deploy appropriate hardware/software solutions at other production sites

WBS 1.5.5.1-Y-12-001 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 are 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 is 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. Tools will be implemented to support visualization and software development and implementation.

In FY09, Y-12 acquired and installed on the cluster two codes in support of material modeling and manufacturing optimization (GULP and ALE3D, respectively) for applications to materials and processes of concern to Y-12. Region-of-interest (ROI) reconstruction code has been implemented and tested successfully using CT projection data acquired on a 420 KeV CT system. This capability allows Y-12 to perform reconstructions of regions of interest acquired by limited field-of-view imagers while maintaining the proper mathematical support. Monte Carlo techniques with parallel resources are being used to study scatter contributions in high-energy (MeV) industrial radiography facility design. Approximately 1600 MCNP simulations have been completed on the cluster during FY09 in support of this project.

Planned activities in FY10:

- Continue evaluations of material modeling and manufacturing optimization codes (GULP and ALE3D, respectively) for applications to materials and processes of concern to Y-12
- Collaborate with KCP on determining physics-based simulation solution needs utilizing SIERRA framework
- Interface with KCP on common welding applications and solution techniques
- Acquire a GPU computing platform and evaluate for Monte Carlo-based transport codes, such as MCNP
- Summarize and review material properties measured at Y-12 for submission to complex-wide materials database
- Investigate the use of UQ techniques in tomographic reconstruction of stockpile components
- Utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Participate in Nuclear Weapons Complex ASC activities

Expected deliverables in FY10:

- Determination of physics-based simulation needs at all production sites

- Demonstrated material modeling code on cluster
- Demonstrated manufacturing optimization code on cluster

Preliminary planned activities in FY11:

- Continue to apply new codes on Y-12 cluster and to utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Continue evaluation of GPU computer as an alternative, general –purpose parallel computing resource
- Participate in Nuclear Weapons Complex ASC activities

WBS 1.5.5.2: User Support Services

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

User Support Services Deliverables for FY10:

- Replacement of Remedy trouble ticket system with Front Range Incident Management module
- Sequoia ID (Dawn) system training to ASC and Alliance partners
- Ongoing customer support metrics
- New training classes and documentation developed for Roadrunner phase III system
- Roll out online unclassified and classified user forum sites
- Reliable and responsive service to users in the ASC tri-lab computing environments.
- Customer support via tri-lab service level agreements for transferring HPC support issues among labs.
- Training for system administrators and users via Web-based short courses
- Online support on classified network
- Partial standup of user support services for remote use of Mesa, commensurate with Mesa capability computing platform system integration readiness milestone and early access to the platform

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.

In FY09, LLNL deployed a new identity/account management tool. LLNL also developed documentation and training materials for the Sequoia Initial Delivery system.

Planned activities in FY10:

- Replace Remedy trouble ticket system with FrontRange Incident Management module
- Deliver the Sequoia ID (Dawn) system training to ASC and Alliance partners
- Evaluate FrontRange Configuration Management module
- Provide ongoing support services for hotline operations, documentation, and training
- Provide input to development of Phase II Identity Management Features

Expected deliverables in FY10:

- Trouble ticket system replacement
- Sequoia ID system training

Preliminary planned activities in FY11:

- Implement configuration management FrontRange module
- Evaluate FrontRange change management, service level management, and problem management modules
- Develop Sequoia system documentation and training materials
- 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.

In FY09, LANL developed training classes and documentation for Hurricane, Turing, and Lobo, and a training class for Roadrunner with focus on efficient hybrid computing. LANL rolled out initial training and documentation Web sites; implemented Request Tracker, a Web-based ticket tracking system and an internal instant message platform; and leveraged various aspects of Request Tracker and the instant message system enabling real-time notifications to on-duty consultants.

Planned activities in FY10:

- Perform ongoing user support for users of ASC/LANL computing resources
- Expand online documentation in both breadth and depth

- Expand communication options with users through user forum collaboration Web sites
- Collaborate with SNL to establish user support procedures for tri-lab users of Mesa

Expected deliverables in FY10:

- Training classes for Roadrunner hybrid computing programming models

Preliminary planned activities in FY11:

- Perform ongoing user support for users of ASC/LANL computing systems
- Continue to improve the quality of support for productive use of computing resources
- Continue to explore alternative communication and support models for providing responsive support within the ASC program

WBS 1.5.5.2-SNL-001 User Support

The User Support team facilitates 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 that 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 trained in HPC rather than application code developers who focus more on application code improvements.

The User 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 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.

In FY09, SNL deployed the knowledge base infrastructure on the classified network, contributed to the Mesa operations planning project, facilitated continuing tri-lab interactions in support of cross-realm authentication and Web-access, and continued hosting the tri-lab user support collaboration meetings.

Planned activities in FY10:

- 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 Mesa (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 user and application support; object is to improve the effectiveness of the principle capacity computing systems in supporting ASC projects
- 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
- Provide support for new ASC platforms/architectures (including Mesa and Sequoia/the Sequoia ID (Dawn))
- 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 FY10:

- Reliable and responsive service to users in the ASC tri-lab computing environments
- Online support on classified network
- Customer support via tri-lab service-level agreement for transferring HPC support issues among labs

Preliminary planned activities in FY11:

- Implement the joint LANL-SNL operational support model for Mesa
- 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 FY10

- Support the ASC PI meeting
- Support for Supercomputing conference and SC09 ASC Research Exhibit
- Host two Predictive Science Panel meetings
- Multimedia and written communications for NA-114

WBS 1.5.5.3-LLNL-001 Program Support

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

In FY09, the most significant change in the Program Support project was the addition of the Sequoia contract.

Planned activities in FY10:

- Continue management of existing Purple and BlueGene/L contracts as well as the new Sequoia contract and its associated D&E contract with IBM for BlueGene/Q
- Manage existing tri-lab contacts and negotiate/execute any new contracts
- Support the ASC PI meeting, the annual Supercomputing conference, Predictive Science Panel meetings, and other meetings and workshops
- Support Predictive Science Academic Alliance Program (PSAAP) collaborations
- Continue closeout of the original Academic Strategic Alliance Program (ASAP) contracts
- Provide support to the ASC Federal program-management office
- Participate in planning with the HEXLABS community (Argonne, LANL, Lawrence Berkeley, LLNL, Oak Ridge, SNL) for a possible future exascale program in partnership with the Office of Science; this may include activities such as workshops
- Plan with HEXLABS community (Argonne, LANL, Lawrence Berkeley, LLNL, Oak Ridge, SNL) for possible future exascale program in partnership with SC (Office of Science); this may include activities such workshops

Expected deliverables in FY10:

- Execution of any new CCE and TLCC-related contracts

Preliminary planned activities in FY11:

- Continue FY10 activities

WBS 1.5.5.3-LLNL-002 Scientific Collaborations

This project provides support for scientific 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.

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 OGAs.

In FY09, LLNL provided project reports directly to the lead collaboration agencies.

Planned activities in FY10:

- 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 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 stress corrosion cracking applications. Detailed project planning and accomplishments are reported separately to DOE through the project PI.
- 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 that the new methods

scale to hundreds of thousands of processors. Detailed project planning and accomplishments are reported separately to DOE through the project PI.

- Continue to collaborate with DTRA Nuclear Weapons Effects Division to investigate Electromagnetic Pulse signature of several new types of Electromagnetic Pulse-Generating Phenomena and begin studies of several new candidates for observed, but unexplained particle jetting to satellite-relevant altitudes.
- Collaborate with the Russian Federation Institutes.

Expected deliverables in FY10:

- Project progress reports to the lead collaboration agencies

Preliminary planned activities in FY11:

- Continuation of the FY10 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.

In FY09, LANL hosted the Predictive Science Panel; provided consultant support to the Federal program efforts to foster collaborations while building support within the predictive science community; supported the PSAAP; and provided support for the Supercomputing conference.

Planned activities in FY10:

- 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 FY10:

- Organization of the ASC tri-lab booth at the SC09 conference
- Predictive Science Panel meeting

Preliminary planned activities in FY11:

- Host the Predictive Science Panel

WBS 1.5.5.3-SNL-001 One Program/Three Labs

The One Program/Three Labs project 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.

In FY09, SNL held the inaugural meeting of the Predictive Engineering Science Panel, an external review of integrated simulation and phenomenological science activities crucial to the establishment of predictive capabilities for engineering assessment and certification of the future stockpile. SNL also contributed to successful exhibits at SC08

and the second Congressional Exhibition on Modeling and Simulation, as well as continued communications and strategic planning support to HQ.

Planned activities in FY10:

- Organize and host second Predictive Engineering Science Panel meeting
- Support PSAAP collaborations
- Provide support to the ASC Federal program office
- Develop and support the SNL and HQ ASC Web sites
- Continue production of high-quality communications materials for HQ and the broader HPC community
- Support the annual Principle Investigator's meetings that expose attendees to technical and programmatic efforts at the three laboratories, and support the Supercomputing Conference and Predictive Science Panel meetings
- Support for 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 collaborations with DTRA, STRATCOM, and the NW Effects User's Group

Expected deliverables in FY10:

- Predictive Engineering Science Panel meeting
- Support for the SC09 Supercomputing Conference

Preliminary planned activities in FY11:

- Host the Predictive Engineering Science Panel
- Support PSAAP collaborations
- Manage 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	Site(s)	Participating Program Offices
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	HQ, LLNL, LANL	NA 121.2 Science Campaigns ASC Campaign
3	Baseline demonstration of UQ aggregation methodology for full-system weapon performance prediction	1	FY12	Q4	HQ, LLNL, LANL, SNL	NA 121.2 Science Campaigns ASC Campaign DSW
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	FY12	Q4	HQ, LLNL, LANL	NA 121.2 Science Campaigns ASC Campaign
5	Assessment of weapon surety status (nuclear safety and physical security) in off-normal transportation scenarios.	1	FY12	TBD	HQ, SNL	NA 121.2
6	Demonstrate predictive capability for weapon system response to short-pulsed neutrons in hostile radiation environment.	1	FY13	TBD	HQ SNL	NA 121.2
7	Full-system safety assessment of damaged weapon immersed in fuel fire for transportation accident scenario.	1	FY14	TBD	HQ, SNL	NA 121.2

Table V-2. Quick Look: Level 2 Milestone Dependencies for FY10[5]

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	Demonstrate ALE-AMR capability	2	FY10	Sep-10	IC	LLNL
TBD	Implement additional physics module sharing in nuclear performance code system	2	FY10	Sep-10	IC	LLNL
TBD	Deliver improved implicit and explicit schemes for both thermal and solid mechanics	2	FY10	Sep-10	IC	LLNL
TBD	Advanced hydrogen EOS delivered for NIF, QMU, and NBI	2	FY10	Sep-10	PEM	LLNL
TBD	Develop and deploy the next-generation multiphase material strength model	2	FY10	Sep-10	PEM	LLNL
TBD	Assess performance of advanced high explosive burn models	2	FY10	Jun-10	PEM	LLNL
TBD	Apply PARADISO and advanced analytic models to assessment of plasma properties such as EOS	2	FY10	Sep-10	PEM	LLNL
TBD	Improved evaluation for the spectrum of neutrons emitted promptly in fission of 239 plutonium	2	FY10	Jun-10	PEM	LLNL
TBD	UQ of primaries with varying marginality	2	FY10	Sep-10	V&V	LLNL
TBD	Secondary computational assessment methodology program (SCAMP) extended to 2D	2	FY10	Sep-10	V&V	LLNL
TBD	Scalable applications preparations and outreach for the Sequoia ID (Dawn)	2	FY10	Jun-10	CSSE	LLNL
TBD	Terascale Simulation Facility 7.5 MW power upgrade; phase 1 of 15 MW upgrade	2	FY10	Sep-10	FOUS	LLNL
TBD	Deploy FrontRange incident management module	2	FY10	Jun-10	FOUS	LLNL
TBD	Certify and accredit Livermore Computing to NAPS 14.2c	2	FY10	Sep-10	FOUS	LLNL
TBD	Enhance code suite physics and optimization in support of the national technical nuclear forensics program and high energy density physics experimental program	2	FY10	Mar-10	IC	LANL
TBD	Enhance code suite physics and optimization in support of directed stockpile work and the national code strategy	2	FY10	Sep-10	IC	LANL

[5] Factors such as FY10 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	Enhance hydro capabilities	2	FY10	Jun-10	IC	LANL
TBD	Enhance code capabilities in support of DSW	2	FY10	Sep-10	IC	LANL
TBD	Improved mix models	2	FY10	Sep-10	PEM	LANL
TBD	High rate improvements to strength models	2	FY10	Sep-10	PEM	LANL
TBD	Advanced reactive flow model for insensitive high explosives	2	FY10	Sep-10	PEM	LANL
TBD	Improved plutonium equation of state	2	FY10	Sep-10	PEM	LANL
TBD	Assessment of primary burn in an ASC code	2	FY10	Sep-10	V&V	LANL
TBD	Thermonuclear applications V&V assessment of physics modeling capabilities in an ASC code	2	FY10	Sep-10	V&V	LANL
TBD	Engineering V&V assessment of a high-explosive assembly stress-state characterization	2	FY10	Sep-10	V&V	LANL
TBD	Roadrunner Phase 3 transition to operational status	2	FY10	Jun-10	CSSE	LANL
TBD	Application enablement on next-generation platforms	2	FY10	Sep-10	CSSE	LANL
TBD	Infrastructure equipment upgrades project	2	FY10	Mar-10	FOUS	LANL
TBD	Coupled thermal and quasistatic failure capabilities for assured safety applications	2	FY10	Sep-10	IC	SNL
TBD	Time-dependent NWM21 radiation environment	2	FY10	Sep-10	IC	SNL
TBD	TRILINOS/SIERRA ToolKit integration	2	FY10	Sep-10	IC	SNL
TBD	Validated physics-based neutron damage effects models for III-V hetero-junction bipolar transistors	2	FY10	Sep-10	PEM	SNL
TBD	Experimentally validated constitutive model for lead-free solder to simulate aging and reliability of solder joints in stockpile components	2	FY10	Sep-10	PEM	SNL
TBD	Conformal decomposition finite element method interface tracking technology for full 3D, parallel, transient capabilities in SIERRA Mechanics	2	FY10	Sep-10	PEM	SNL
TBD	Abnormal thermal safety response QMU baseline for B61-3/4/10	2	FY10	Sep-10	V&V	SNL
TBD	Deliver a prediction, with bounded uncertainties, of the response of a representative AF&F sub-circuit under a combined neutron/gamma hostile environment	2	FY10	Jun-10	V&V	SNL

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	SIERRA code feature and capability test coverage assessment	2	FY10	Sep-10	V&V	SNL
TBD	Evaluate advanced memory subsystems	2	FY10	Sep-10	CSSE	SNL
TBD	Deployment of a common capacity computing environment	2	FY10	Sep-10	CSSE	LLNL, LANL, SNL
TBD	Visualization-on-platform technology	2	FY10	Sep-10	CSSE	LANL, SNL

Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY11

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	Investigate boost physics issues	2	FY11	Sep-11	IC	LLNL
TBD	Implement algorithms for improved loading on structures	2	FY11	Sep-11	IC	LLNL
TBD	Multiphase Pu EOS variations delivered for QMU	2	FY11	Sep-11	PEM	LLNL
TBD	Assess sensitivity of material properties models to advanced high-explosive burn models	2	FY11	Jun-11	PEM	LLNL
TBD	Apply new simulation and modeling capability	2	FY11	Sep-11	PEM	LLNL
TBD	Scalable applications preparations and outreach for Sequoia	2	FY11	Jun-11	CSSE	LLNL
TBD	Data analysis plan for Sequoia	2	FY11	Jun-11	CSSE	LLNL
TBD	Terascale Simulation Facility 7.5 MW power upgrade; phase 2 of 15 MW upgrade	2	FY11	Sep-11	FOUS	LLNL
TBD	Enhance the code suite physics in support of the TBI and the Predictive Capabilities Framework—initial conditions for boost I	2	FY11	Mar-11	IC	LANL
TBD	2D Primary design capability in support of DSW	2	FY11	Jun-11	IC	LANL
TBD	Enhance the code suite physics in support of the national code strategy and robustness through improved testing	2	FY11	Sep-11	IC	LANL
TBD	Explore algorithmic diversity in transport methods	2	FY11	Sep-11	IC	LANL
TBD	Explore algorithmic diversity in transport methods	2	FY11	Sep-11	IC	LANL
TBD	Initial implementation of in-line NLTE capability	2	FY11	Mar-11	PEM	LANL
TBD	Subgrid damage model for fragmentation problems	2	FY11	Sep-11	PEM	LANL
TBD	Next generation of charged-particle data capabilities	2	FY11	Jun-11	PEM	LANL
TBD	Verification toolbox	2	FY11	Sep-11	V&V	LANL
TBD	Benchmark evaluation of predictive capability for boost using LANL Boost Validation Suite (BVS)	2	FY11	Sep-11	V&V	LANL
TBD	Engineering V&V assessment of a simplified weapon subassembly subjected to shock loading	2	FY11	Sep-11	V&V	LANL
TBD	VIZ cluster upgrade project	2	FY11	Sep-11	CSSE	LANL
TBD	Cavity SGEMP predictive capability for realistic RB geometry	2	FY11	Sep-11	IC	SNL

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/ Subprogram(s)	Site(s)
TBD	One-way coupling of re-entry aerodynamics with ablation to structural analysis simulations	2	FY11	Sep-11	IC	SNL
TBD	Coupled blast-to-structure simulation capability supporting NW safety and security	2	FY11	Sep-11	IC	SNL
TBD	Coupled meshing to SIERRA capability	2	FY11	Sep-11	IC	SNL
TBD	Mechanical material failure models for B61 laydown environments	2	FY11	Sep-11	PEM	SNL
TBD	Next-generation NT simulation arc physics models	2	FY11	Sep-11	PEM	SNL
TBD	Gas transfer system performance models for hydrogen embrittlement	2	FY11	Sep-11	PEM	SNL
TBD	W80 abnormal mechanical QMU	2	FY11	Sep-11	V&V	SNL
TBD	Computational UQ for the QASPR	2	FY11	Sep-11	V&V	SNL
TBD	Cavity SGEMP V&V	2	FY11	Dec-10	V&V	SNL
TBD	Aerodynamic performance V&V study	2	FY11	Sep-11	V&V	SNL
TBD	Deliver visualization-based analysis tools in support of FY12 full system safety milestone	2	FY11	Sep-11	CSSE	SNL
TBD	Mesa Capability Computing Platform integration readiness	2	FY10	Dec-10	CSSE	LANL, SNL
TBD	Mesa Capability Computing Platform production readiness	2	FY11	Jun-11	CSSE	LANL, SNL
TBD	Develop TOSS 2.0	2	FY11	Sep-11	CSSE	LLNL, LANL, SNL
TBD	Deploy TLCC11 clusters	2	FY11	Sep-11	FOUS	LLNL, LANL, SNL

Detailed Milestone Descriptions for FY10

Milestone (ID#): Scalable applications preparations and outreach for the Sequoia ID (Dawn)		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: June 30, 2010		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The SAP effort will develop the knowledge base, documentation, and training to provide ASC code teams with support for utilization of the Sequoia ID (Dawn). SAP will actively engage tri-lab code teams to address their needs in porting codes to the Sequoia ID (Dawn), exploring options for multi-core utilization, characterizing performance issues for the codes. For FY10, tri-lab code teams will be surveyed for needs. One or more multi-physics codes will be engaged to characterize the Sequoia ID (Dawn) performance, analyze bottlenecks and load balance issues, and to develop strategies for improving performance targeting the Sequoia system.		
Completion Criteria: A report covering the performance findings, and recommended techniques and strategies for the codes studied will be prepared.		
Customer: ASC Integrated Codes Program		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, will be prepared as a record of milestone completion. A "handoff" letter accompanying the report about the project findings and recommendations to the IC program will be document.		
Supporting Resources: ASC the Sequoia ID (Dawn) system		
Supporting Milestones:		

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

Codes/Simulation Tools Employed: The Sequoia ID (Dawn) software environment and IBM Sequoia simulator				
Contribution to the ASC Program: More effective use of Sequoia computer system				
Contribution to Stockpile Stewardship: More effective use of Sequoia computer system				
No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	No access to IBM Sequoia simulator	Low	Medium	Low

<table>
<tr><td colspan="6">Milestone (ID#): Terascale Simulation Facility 7.5 MW power upgrade; phase 1 of 15-MW upgrade</td></tr>
<tr><td colspan="2">Level: 2</td><td colspan="2">Fiscal Year: FY10</td><td colspan="2">DOE Area/Campaign: ASC</td></tr>
<tr><td colspan="6">Completion Date: September 30, 2010</td></tr>
<tr><td colspan="6">ASC nWBS Subprogram: FOUS</td></tr>
<tr><td colspan="6">Participating Sites: LLNL</td></tr>
<tr><td colspan="6">Participating Programs/Campaigns: ASC</td></tr>
<tr><td colspan="6">Description: Complete distribution of 7.5 MW computational power upgrade to the TSF West Computer Room.</td></tr>
<tr><td colspan="6">Completion Criteria: This project is complete when LC has received certification of acceptance testing completion from the LLNL electrical readiness contractor Electro Links in conjunction with LLNL and the Lab's Labor only organization, GSE.</td></tr>
<tr><td colspan="6">Customer: Weapons Complex Integration</td></tr>
<tr><td colspan="6">Milestone Certification Method:
A report will be written as a record of milestone completion.
A "handoff" memo will be generated and will be addressed to Bruce Goodwin, Weapons Complex Integration Associate Director.</td></tr>
<tr><td colspan="6">Supporting Resources: GSE Labor Force</td></tr>
<tr><td colspan="6">Supporting Milestones:</td></tr>
<tr><td colspan="2">Program</td><td colspan="2">Title</td><td colspan="2">Due Date</td></tr>
<tr><td colspan="2">N/A</td><td colspan="2">N/A</td><td colspan="2">N/A</td></tr>
<tr><td colspan="6">Codes/Simulation Tools Employed: None</td></tr>
<tr><td colspan="6">Contribution to the ASC Program: Is required for Sequoia power requirements. Sequoia will be used as an UQ studies platform, and will be integrated in 2012.</td></tr>
<tr><td colspan="6">Contribution to Stockpile Stewardship:</td></tr>
<tr><td rowspan="2">No.</td><td rowspan="2">Risk Description</td><td colspan="4">Risk Assessment (low, medium, high)</td></tr>
<tr><td>Consequence</td><td>Likelihood</td><td colspan="2">Exposure</td></tr>
<tr><td>1</td><td>Unable to site Sequoia because project is not completed</td><td>High</td><td>Low</td><td colspan="2">Medium</td></tr>
</table>

Milestone (ID#): Deploy FrontRange incident management module		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: June 30, 2010		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: LC's current trouble tracking system, Remedy, will be replaced with the FrontRange IT Service Management product. FrontRange is a product based on industry best practices, the Information Technology Infrastructure Library (ITIL) processes. We will use FrontRange's incident management module to provide core service desk functionality to efficiently identify, respond to, and track issues and service requests needing resolution. This milestone is complete when all LC incidents are managed via the FrontRange system and service level agreement metrics reports are available.		
Completion Criteria: This project is complete when LC's daily tickets are entered and tracked via the FrontRange Incident Management module and the Remedy system is being used solely for historical reporting purposes.		
Customer: LC		
Milestone Certification Method: A report will be written that documents completion of the FrontRange incident management module for problem tracking and reporting at LC. A memo will be written that documents the handoff of the developed capability to Michel McCoy, ASC Program Leader at LLNL.		
Supporting Resources: LLNL Institutional Services Division, CSSE personnel		
Supporting Milestones:		

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

Codes/Simulation Tools Employed: None
Contribution to the ASC Program: Required for successful management of ASC computer system resources and timely resolution of customer issues.
Contribution to Stockpile Stewardship:

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	Trouble ticket system is not integrated in a timely fashion due to unforeseen software issues	Low	Medium	Low

Milestone (ID#): Certify and accredit Livermore Computing to NAPS 14.2c		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: September 30, 2010		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: LC's cyber security plan accreditation expires March 31, 2010. A new plan must be written to be compliant with the NNSA's NAPS 14.2c cyber security policy. This new policy requires a complete overhaul of LLNL's classified security plans. LC will be a major contributor in the effort to create a Site Security Component Library, which will be used as a basis for LC's NAPS compliant plan. All new security test and evaluation plans will be developed to demonstrate NAPS compliance. Ultimately, the LC Information System Security Plan (ISSP) will be developed and delivered. This milestone is complete when the ISSP and the security test and evaluation plans have been submitted and LC has received interim approval to test.		
Completion Criteria: This milestone is complete when the ISSP and the security test and evaluation plans have been submitted and LC has received interim approval to test.		
Customer: LLNL ASC Program		
Milestone Certification Method: A report will be prepared to document completion of this milestone. A memo will be supplied that documents the Livermore Site Office's approval for LC to execute its Security Test and Evaluation plans.		
Supporting Resources: Livermore Site Office, LLNL Computer Security Program, LLNL's iSRD Team		
Supporting Milestones:		

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

Codes/Simulation Tools Employed: None
Contribution to the ASC Program: Approval of this plan is required to continue operating the LC, which contains all the ASC HPC resources.
Contribution to Stockpile Stewardship:

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	Security Plan is not approved on time.	High	Medium	High

Milestone (ID#): Roadrunner Phase 3 transition to operational status		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: June 30, 2010		
ASC nWBS Subprogram: CSSE		
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.		
Completion Criteria: Roadrunner system running with a select set of users; complete infrastructure support provided, including I/O, networking, and archiving; advanced Architecture operation state software stack is provided; system running qualified hybrid applications; and ASC and weapons science applications target for Roadrunner have demonstrated operational readiness.		
Customer: Selected set of users requiring hybrid applications system.		
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: System and application analyst support.		
Supporting Milestones:		
Program	**Title**	**Due Date**
N/A	N/A	N/A
Codes/Simulation Tools Employed: Hybrid programming tools as well as traditional parallel programming tools as specified in the statement of work.		
Contribution to the ASC Program: The hybrid Roadrunner system provides leadership opportunity in code design and development for future capability systems such as multi-core and hybrid. Brings compute power for ASC Program into the petaFLOPS range, enabling scientific calculations in support of ASC Roadmap.		
Contribution to Stockpile Stewardship: Provides crucial compute cycles and performance to weapons simulations and physics codes for taking the predictive capability to the next level.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	**Likelihood**	**Exposure**
1.	Weapons application not ready to run by 1QFY10	High	Low	Medium

Milestone (ID#): Application enablement on next-generation platforms		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: September 30, 2010		
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, one emphasis will be on development of tools for memory, communications and performance tracking. This work will focus on a radiation-hydrodynamics workload from the ASC Eulerian code project. The milestone will provide the ASC community with memory and communication tools that aid in the understanding of code performance on current generation advanced platforms (Roadrunner) and on future systems. Combining the findings from the data movement analysis with architectural trends, we will help focus the application architecture for the next generation capability machines at the LANL. Specifically, we will evaluate the potential performance gains in the radiation-hydrodynamics application under consideration prior to a potential re-engineering of the code for heterogeneous architectures.		
Completion Criteria: Results documented as a report. Tools made available to the ASC community.		
Customer: Integrated Codes		
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: Access to Roadrunner and other supercomputer resources		

Supporting Milestones:

Program	Title	Due Date
ASC/CSSE	Mesa capability computing platform production capability readiness	Jun-11
ASC/CSSE	Roadrunner Phase 3 transition to operational status	Jun-10

Codes/Simulation Tools Employed: ASC Eulerian code project
Contribution to the ASC Program: The milestone will provide the ASC community with memory and communication tools that aid in the understanding of code performance on current generation advanced platforms (Roadrunner) and on future systems.
Contribution to Stockpile Stewardship: Combining the findings from the data movement analysis with architectural trends, we will help focus the application architecture for the next generation capability machines.

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	**Likelihood**	**Exposure**
1	Loss of key personnel	High	Medium	Medium

Milestone (ID#): Infrastructure equipment upgrades project		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: March 31, 2010		
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 consist 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 Mesa in 2010.		
Completion Criteria: Additional 7.2MW and chilling capacity completed. All electrical and chilling equipment installed and operational. SCC ready for Mesa site prep work.		
Customer: ASC, ACES (Mesa capability computing platform)		
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: FOUS funding		
Supporting Milestones:		

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

Codes/Simulation Tools Employed: N/A
Contribution to the ASC Program: Provides infrastructure for Mesa.
Contribution to Stockpile Stewardship: Provides power for Mesa that provides CPU cycles for SSP.

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1.	Installation schedule not met.	High	Low	Medium

Milestone (ID#): Evaluate advanced memory subsystems		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: September 30, 2010		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Develop a next-generation memory system architecture to increase performance of SNL applications in partnership with industry and academia. SNL is providing simulation capabilities, performing architectural analysis, and supplying application expertise. Simulation of this system shall show order of magnitude improvements in memory system performance. Simulation shall also be used to design new memory access protocols, evaluate the power/performance tradeoffs of different memory system topologies, and determine how overall system balance (processor core count and capabilities, interconnect bandwidths, and memory hierarchy sizes) should be altered to take advantage of the new memory system. Additionally, the performance impact of "Smart" memory operations including atomic memory operations, flexible data movement offload, and in-memory synchronization, shall be quantified in simulation and in an FPGA prototype. The architecture shall be defined with an industrial partner who will develop a concept prototype.		
Completion Criteria: Completion of program review and final document published as a SAND report.		
Customer: NNSA/ASC HQ, CSSE program managers and platform design team members, IC program managers and algorithm/application developers.		
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: CSSE Simulation & Tools program, CSSE Advanced Systems program		
Supporting Milestones:		

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

Codes/Simulation Tools Employed: This milestone leverages CSSE SST macro and micro simulation tools, CSSE performance analysis and modeling tools in addition to commercial development tools, such as the Xilinx FPGA tool suite.
Contribution to the ASC Program: Future platform architectures will contain advanced memory technologies that will provide significantly enhanced performance for a large fraction of ASC codes.
Contribution to Stockpile Stewardship: Platforms

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	**Likelihood**	**Exposure**
1	Insufficient participation from partners	High	Low	Med

<table>
<tr><td colspan="6">Milestone (ID#): Deployment of a common capacity computing environment</td></tr>
<tr><td colspan="2">Level: 2</td><td colspan="2">Fiscal Year: FY10</td><td colspan="2">DOE Area/Campaign: ASC</td></tr>
<tr><td colspan="6">Completion Date: September 30, 2010</td></tr>
<tr><td colspan="6">ASC nWBS Subprogram: CSSE</td></tr>
<tr><td colspan="6">Participating Sites: LLNL, LANL, SNL</td></tr>
<tr><td colspan="6">Participating Programs/Campaigns: ASC</td></tr>
<tr><td colspan="6">Description: Deploy additional CCE capabilities for capacity computing environment, working towards a responsive and more efficient infrastructure to support computing for QMU and predictivity.</td></tr>
<tr><td colspan="6">Completion Criteria: Deploy CCE capabilities developed during FY09 including: major upgrades to the common operating system (TOSS); O|SS; Workload Characterization; Application Monitoring; and Gazebo Test and Analysis Suite. Demonstrate the tri-lab CCE software stack on the production ASC TLCC systems. Develop new capabilities in the continuing FY10 CCE projects. The tri-labs will continue to do gap and risk analysis of the CCE software stack and add new projects, as needed, to address high-priority gaps.</td></tr>
<tr><td colspan="6">Customer: Users of the ASC capacity systems</td></tr>
<tr><td colspan="6">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="6">Supporting Resources: ASC CSSE sub-program</td></tr>
<tr><td colspan="6">Supporting Milestones:
<table>
<tr><th>Program</th><th>Title</th><th>Due Date</th></tr>
<tr><td>N/A</td><td>N/A</td><td>N/A</td></tr>
</table></td></tr>
<tr><td colspan="6">Codes/Simulation Tools Employed: N/A</td></tr>
<tr><td colspan="6">Contribution to the ASC Program: Common Computing Environment for capacity computing</td></tr>
<tr><td colspan="6">Contribution to Stockpile Stewardship: Easy-to-use capacity computing resource</td></tr>
<tr><td rowspan="2">No.</td><td rowspan="2">Risk Description</td><td colspan="4">Risk Assessment (low, medium, high)</td></tr>
<tr><td>Consequence</td><td colspan="2">Likelihood</td><td>Exposure</td></tr>
<tr><td></td><td>None</td><td></td><td colspan="2"></td><td></td></tr>
</table>

Milestone (ID#): Visualization-on-platform technology		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: September 2010		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL/SNL		
Participating Programs/Campaigns: ASC		
Description: Visualization and analysis of petascale data is limited by several factors which must be addressed as ACES delivers Mesa. Two primary difficulties are: 1) Performance of interactive rendering, which is the most computationally intensive portion of the visualization process. For terascale platforms, commodity clusters with GPUs have been used for interactive rendering. For petascale platforms, visualization and rendering may be able to run efficiently on the supercomputer platform itself; and 2) I/O bandwidth, which limits how much information can be written to disk. If we simply analyze the sparse information that is saved to disk, we miss the opportunity to analyze the rich information produced every time step by the simulation. For the first issue, we are pursuing in-situ analysis, in which simulations are coupled directly with analysis libraries at runtime. This milestone will evaluate the visualization and rendering performance of current and next-generation supercomputers in contrast to GPU-based visualization clusters, and evaluate the performance of common analysis libraries coupled with the simulation that analyze and write data to disk during a running simulation. This milestone will explore, evaluate and advance the maturity level of these technologies and their applicability to problems of interest to the ASC program.		
Completion Criteria: Results documented as a report. Tools made available to the ASC community.		
Customer: ASC		
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: Access to Roadrunner and other supercomputer resources.		
Supporting Milestones:		
Program	**Title**	**Due Date**
ASC/CSSE	Mesa capability computing platform production capability readiness	Jun-11
ASC/CSSE	VIZ cluster upgrade project	Sept-11
Codes/Simulation Tools Employed: Various research and prototype visualization codes		
Contribution to the ASC Program: Evaluates potential future systems which could save cost and/or improve workflow efficiency		
Contribution to Stockpile Stewardship: Evaluates potential future systems which could save cost and/or improve workflow efficiency		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	**Likelihood**	**Exposure**
1	Loss of key personnel	High	Medium	Medium
2	Delays in Mesa schedule	Medium	Medium	Medium

Milestone Descriptions for Preliminary FY11

Milestone (ID#): Scalable applications preparations and outreach for Sequoia		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: June 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Building on the FY10 milestone, the SAP effort will extend the knowledge base, documentation, and training for the Sequoia ID (Dawn). SAP will engage tri-lab code teams to address their needs in porting codes to the Sequoia ID (Dawn), and preparing for arrival of Sequoia in FY12. For FY11, additional multi-physics codes will be engaged to characterize the Sequoia ID (Dawn) performance, analyze bottlenecks and load balance issues, and to develop strategies for improving performance targeting the Sequoia system. Additional refinements to understanding and performance enhancements will be achieved for the codes previously targeted.		

Milestone (ID#): Data analysis plan for Sequoia		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: June 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This plan defines LLNL's site strategy for Tri-lab large-scale data analysis on Sequoia and its related infrastructure environment. It is a Sequoia-specific follow-on to the FY08 Level 2 Milestone "Infrastructure Plan for ASC Petascale Environments" (LLNL Technical Report 402112), and addresses Phase 2 (i.e., FY11-FY13) petascale data analysis on Sequoia. The plan will provide detailed description of tools and capabilities for data movement, data analysis, visualization, and mass store archiving that will be deployed to support petascale data tasks for Sequoia users. While this plan primarily will describe LLNL site-specific development, deployment and operational support for Sequoia, the objective is for the described data capabilities to be used by the broad Tri-lab user community in support of Sequoia UQ and weapon science simulations.		

Milestone (ID#): Terascale Simulation Facility 7.5 MW power upgrade; phase 2 of 15-MW upgrade		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: 4QFY11		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Complete distribution of 7.5 MW computational power upgrade to the TSF East Computer Room.		

Milestone (ID#): VIZ cluster upgrade project		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: Sep-11		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: The focus of this milestone will be to provide the necessary resources to visualize output generated on petascale clusters. Requirements will be developed, equipment will be purchased, and the cluster will be integrated into the LANL computer center. The milestone will be complete when the cluster is ready for production work.		

Milestone (ID#): Deliver visualization-based analysis tools in support of FY12 full system safety milestone		
Level: 2	**Fiscal Year**: FY10	**DOE Area/Campaign**: ASC
Completion Date: September 30, 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: The Full System Safety Milestone requires analysis of complex ensembles of runs, but current tools do not support rich analysis of this data, Instead, analysts rely on summarized analysis of the data. SNL's Data Analysis and Visualization team, in partnership with the SNL V&V program, are targeting ensemble analysis tools to solve analysis problems in support of their FY12 Full System Safety Milestone. It is critical to have ensemble analysis tools in place in FY10, so that those tools can directly support the analysis needed to accomplish the V&V milestone that will occur in FY11 and FY12. A PCMM evaluation shall be made for this capability.		

Milestone (ID#): Mesa Capability Computing Platform system integration readiness		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: December 30, 2010		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: Mesa is ready for integration into the LANL computing center. System hardware delivery to LANL has been completed. The system installation at LANL has been completed. System software for the system has been delivered, tested, and demonstrated. On-site capability scaling testing has been completed. Mesa is ready for on-site integration into the local and remote computing infrastructure, including the user software environment.		
Completion Criteria: Follows the ASC Level 2 Milestone criteria for capability platforms: system hardware deliveries from vendor to site are complete, including the basic hardware to integrate "the system" as contractually defined; installation of the system by the contractor on-site to the extent that is contractually required is substantially complete; hardware acceptance testing as contractually required is completed; in general, contractual requirements for formal hardware acceptance have been substantially completed; system software needed for basic operation of the system is delivered, tested, and demonstrated to be operational; vendor has completed on-site capability scaling testing and demonstration; and system is ready to begin on-site integration into local computing environment.		
Customer: NNSA/ASC HQ, tri-lab ASC program managers responsible for CCCs, SSP, tri-lab weapons applications community.		
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: CSSE, FOUS, Platform funding, ACES program managers, LANL facilities.		

Supporting Milestones:

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

Codes/Simulation Tools Employed: N/A
Contribution to the ASC Program: Provides production capability compute cycles to ASC Program including scalable performance.
Contribution to Stockpile Stewardship: Primary production capability platform for the ASC Program.

No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1.	Contract not awarded on schedule	High	Low	Medium

Milestone (ID#): Mesa Capability Computing Platform production readiness		
Level: 2	**Fiscal Year**: FY11	**DOE Area/Campaign**: ASC
Completion Date: Jun-11		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL, LANL		
Participating Programs/Campaigns: ASC		
Description: Mesa shall achieve Production Capability Readiness as defined by the Capability Platform Level 2 Milestones Working Group. In summary, this includes the platform is made available for capability-class projects; all system software, tools, utilities and user support processes are available and fully functional; ASC applications targeted for the platform are ported and made available to designers, analysts, and engineers; the platform has demonstrated acceptable reliability performance targets.		

Milestone (ID#): Develop TOSS 2.0	
Fiscal Year: FY11	**DOE Area/Campaign**: ASC
Completion Date: Sept. 2011	
ASC nWBS Subprogram: CSSE	
Participating Sites: LLNL, SNL, LANL	
Participating Programs/Campaigns: ASC	
Description: Deploy CCE capabilities developed during FY10 and FY11, including the next major release of common operating system (TOSS 2.0) and software stack. Prepare for deployment of the next generation of the ASC TLCC systems, which may include hardware and software integration and testing for the tri-lab environment.	

Milestone (ID#): Deploy TLCC11 clusters	
Fiscal Year: FY11	**DOE Area/Campaign**: ASC
Completion Date: Sept. 2011	
ASC nWBS Subprogram: FOUS	
Participating Sites: LLNL, SNL, LANL	
Participating Programs/Campaigns: ASC	
Description: Deploy next generation ASC TLCC systems into production computing environments.	

VI. ASC Roadmap Drivers for FY10–FY11

Table VI-1. ASC Roadmap Drivers for FY10-11[6]

Focus Area 2: ESTABLISH A VALIDATED PREDICTIVE CAPABILITY FOR KEY PHYSICAL PHENOMENA	
2010	Science-based models for neutron tube simulations

[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 (R) and Targets (T)

Performance Indicators	FY 2005 Results	FY 2006 Results	FY 2007 Results	FY 2008	FY 2009	FY 2010	FY 2011	FY 2012	FY 2013	FY 2014	Endpoint Target
Strategic Goal 2.1 (Nuclear Deterrent) GPRA Unit Program Goal 2.1.30.00, Advanced Simulation and Computing Campaign											
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)[a]	N/A	R: 50%	R: 63% T : 63%	R: 72% T: 72%	T: 80%	T: 85%	T: 90%	T: 95%	T: 100%	N/A	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 Security Enterprise. This measure is meant to show how quickly ASC codes are being adopted by the user community in place of legacy codes.
The cumulative percentage reduction in the use of calibration "knobs" to successfully simulate nuclear weapons performance (Long-term Outcome)[a]	N/A	R: 2%	R: 8% T : 8%	R: 16% T: 16%	T: 25%	T: 30%	T: 35%	T: 40%	T: 45%	T: 50%	By 2024, 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.
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 resolution (Long-term Outcome)[a]	N/A	R: 10%	R: 25% T : 25%	R: 37% T: 37%	T: 50%	T: 60%	T: 65%	T: 70%	T: 80%	T:85%	By 2015, 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.
The cumulative percentage of simulation turnaround time reduced while using modern ASC codes (Efficiency)[a]	N/A	R: 6%	R: 7% T : 7%	R: 13% T: 13%	T: 13%	T: 15%	T: 20%	T: 27%	T: 34%	T: 42%	By 2015, 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.

[a] Performance measures were revised in 2007 to be consistent with new program roadmap.

IX. Budget

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.1 Integrated Codes					
1.5.1.1 **Modern Multi-physics Codes**		49.261	0.000	0.000	49.261
	LLNL	29.909			29.909
	LANL	19.352			19.352
	SNL				0.000
	Other				0.000
1.5.1.2 **Legacy Codes**		1.793	0.000	0.000	1.793
	LLNL	0.650			0.650
	LANL	1.143			1.143
	SNL				0.000
	Other				0.000
1.5.1.3 **Engineering Codes**		26.925	0.000	0.000	26.925
	LLNL	1.900			1.900
	LANL				0.000
	SNL	25.025			25.025
	Other				0.000
1.5.1.4 **Focused Research, Innovation & Collaboration**		30.193	0.000	19.800	49.993
	LLNL	11.534		1.500	13.034
	LANL	9.109			9.109
	SNL	9.550			9.550
	Other			18.300	18.300
1.5.1.5 **Emerging & Specialized Codes**		13.910	0.000	0.000	13.910
	LLNL	4.700			4.700
	LANL	6.310			6.310
	SNL	2.900			2.900
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.2 Physics and Engineering Models					
1.5.2.1 **Theoretical Models and Experimental Integration**		23.567	0.000	0.000	23.567
	LLNL	7.132			7.132
	LANL	11.735			11.735
	SNL	4.700			4.700
	Other				0.000
1.5.2.2 **Model Implementation**		22.292	0.000	0.000	22.292
	LLNL	2.524			2.524
	LANL	10.768			10.768
	SNL	9.000			9.000
	Other				0.000
1.5.2.3 **Fundamental Physics Codes and Application**		9.430	0.000	0.000	9.430
	LLNL	4.730			4.730
	LANL				0.000
	SNL	4.700			4.700
	Other				0.000
1.5.2.4 **Material Data Libraries**		3.390	0.000	0.000	3.390
	LLNL	1.187			1.187
	LANL	2.203			2.203
	SNL				0.000
	Other				0.000
1.5.2.5 **Russian Programs**		1.510	0.000	0.000	1.510
	LLNL	0.500			0.500
	LANL	0.510			0.510
	SNL	0.500			0.500
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.3 Verification and Validation					
1.5.3.1 **V&V Methods**		15.250	0.000	0.000	15.250
	LLNL	3.250			3.250
	LANL	7.400			7.400
	SNL	4.600			4.600
	Other				0.000
1.5.3.2 **Primary V&V Assessments**		5.149	0.000	0.000	5.149
	LLNL	2.260			2.260
	LANL	2.889			2.889
	SNL				0.000
	Other				0.000
1.5.3.3 **Secondary V&V Assessments**		4.633	0.000	0.000	4.633
	LLNL	2.105			2.105
	LANL	2.528			2.528
	SNL				0.000
	Other				0.000
1.5.3.4 **Engineering V&V Assessments**		14.294	0.000	0.000	14.294
	LLNL	0.400			0.400
	LANL	1.994			1.994
	SNL	11.900			11.900
	Other				0.000
1.5.3.5 **Specialized V&V Assessments**		3.200	0.000	0.000	3.200
	LLNL	1.100			1.100
	LANL				0.000
	SNL	2.100			2.100
	Other				0.000
1.5.3.6 **Data Validation & Archiving**		2.081	0.000	0.000	2.081
	LLNL	1.033			1.033
	LANL	0.748			0.748
	SNL	0.300			0.300
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.4 Computational Systems and Software Environment					
1.5.4.1 **Capability Systems**		1.899	45.900	0.000	47.799
	LLNL	0.989	1.040		2.029
	LANL	0.910	42.360		43.270
	SNL		2.500		2.500
	Other				0.000
1.5.4.2 **Capacity Systems**		0.000	0.000	0.000	0.000
	LLNL				0.000
	LANL				0.000
	SNL				0.000
	Other				0.000
1.5.4.3 **Advanced Systems**		4.012	25.966	6.000	35.978
	LLNL	1.102	22.366	3.000	26.468
	LANL		3.600		3.600
	SNL	2.910		3.000	5.910
	Other				0.000
1.5.4.4 **System Software and Tools**		21.669	1.000	0.000	22.669
	LLNL	6.903			6.903
	LANL	7.766	1.000		8.766
	SNL	7.000			7.000
	Other				0.000
1.5.4.5 **Input/Output, Storage Systems, and Networking**		15.367	5.000	0.200	20.567
	LLNL	9.500			9.500
	LANL	2.527	2.500	0.100	5.127
	SNL	3.340	2.500	0.100	5.940
	Other				0.000
1.5.4.6 **Post-processing Environments**		11.419	1.000	0.000	12.419
	LLNL	4.200			4.200
	LANL	3.594	1.000		4.594
	SNL	3.625			3.625
	Other				0.000
1.5.4.7 **Common Computing Environment**		8.865	0.000	0.615	9.480
	LLNL	3.675		0.615	4.290
	LANL	2.660			2.660
	SNL	2.530			2.530
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.5 Facility Operations and User Support					
1.5.5.1 **Facilities, Operations and Communications**		36.078	84.000	1.000	121.078
	LLNL	12.669	43.000		55.669
	LANL	17.509	33.000		50.509
	SNL	5.900	8.000		13.900
	Other			1.000	1.000
1.5.5.2 **User Support Services**		8.654	0.000	1.040	9.694
	LLNL	3.331		1.040	4.371
	LANL	3.223			3.223
	SNL	2.100			2.100
	Other				0.000
1.5.5.3 **Collaborations**		20.731	0.000	9.032	29.763
	LLNL	1.897		0.050	1.947
	LANL	0.843		0.439	1.282
	SNL	1.217		1.005	2.222
	Other	16.774		7.538	24.312

	PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
Integrated Codes	122.082	0.000	19.800	141.882
Physics and Engineering Models	60.189	0.000	0.000	60.189
Verification and Validation	44.607	0.000	0.000	44.607
Computational Systems and Software Engineering	63.231	78.866	6.815	148.912
Facility Operations and User Support	65.463	84.000	11.072	160.535
Total	355.572	162.866	37.687	556.125
	64%	29%	7%	556.125

	PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
Integrated Codes				141.882
Physics and Engineering Models				60.189
Verification and Validation				44.607
Computational Systems and Software Engineering				148.912
Facility Operations and User Support				160.535
				556.125
	355.572	162.866	37.687	556.125

	PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
LLNL	119.180	66.406	6.205	191.791
LANL	115.721	83.460	0.539	199.720
SNL	103.897	13.000	4.105	121.002
Other	16.774	0.000	26.838	43.612
	355.572	162.866	37.687	556.125

Month	Monthly Cost	Cumulative Cost
Oct	36.799	36.799
Nov	29.533	66.332
Dec	36.175	102.507
Jan	29.934	132.441
Feb	33.503	165.944
Mar	47.381	213.325
Apr	44.777	258.102
May	51.951	310.053
Jun	48.769	358.822
Jul	38.183	397.005
Aug	43.517	440.522
Sep	65.858	506.380

Appendix A. Glossary

1D	One Dimensional
2D	Two Dimensional
3D	Three Dimensional
ACES	Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AMR	Adaptive Mesh Refinement
ASAP	Academic Strategic Alliance Program
ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
CAD	Computer Aided Design
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CRASH	Center for Radiative Shock Hydrodynamics
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
D&E	Development and Engineering
DoD	Department of Defense
DOE	Department of Energy
DSL	Domain-Specific Language
DSW	Directed Stockpile Work
DTRA	Defense Threat Reduction Agency
EOS	Equation of State
FPGA	Field Programmable Gate Arrays
FOUS	Facility Operations and User Support (WBS 1.5.5)
FVM	Finite Volume Method
GA	General Availability
GB	Gigabytes
GB/s.	Gigabytes per Second
GPFS	Global Parallel File System
HTGL	High Temperature Gasdynamics Laboratory

GPU	Graphical Processing Units
HPC	High-Performance Computing
HPSS	High-Performance Storage System
HQ	ASC Headquarters
I/O	Input/Output
IAA	Institute for Advanced Architectures and Algorithms
ID	Initial Delivery
JEI	Joint Exascale Initiative
KCP	Kansas City Plant
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LC	Livermore Computing Center
LDAP	Lightweight Directory Access Protocol
LDCC	Laboratory Data Communications Center
LEP	Life Extension Program
LES	Large-Eddy Simulation
LLNL	Lawrence Livermore National Laboratory
LWK	Lightweight Kernel
MEMS	Microelectromechanical Systems
MPI	Message Passing Interface
NAPS	NNSA Policy Letters (known as…)
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
OGA	Other Government Agency
PECOS	Center for Predictive Engineering and Computational Sciences
PLFS	Parallel Log Structured File System
PRISM	Prediction of Reliability, Integrity, and Survivability of Microsystems (Center for…)
PSAAP	Predictive Science Academic Alliance Program

PSI	Parallel Storage Interface
QMU	Quantification of Margins and Uncertainties
R&D	Research and Development
RANS	Reynolds-Averaged Navier Stokes
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
SAN	Storage Area Network
SAP	Scalable Applications Preparations
SCC	Nicholas C. Metropolis Center for Modeling and Simulation
SCF	Secure Computing Facility
SFI	Significant Finding Investigation
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
SU	Scalable Unit(s)
TBI	Thermonuclear Burn Initiative
TLCC	Tri-Lab Linux Capacity Cluster
TLCC07	Tri-Lab Linux Capacity Cluster for 2007
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
UQ	Uncertainty Quantification
V&V	Verification and Validation
VTF	Virtual Test Facility
WAN	Wide Area Network
WC Tool	Workload Characterization Tool

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 Manuel 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 Manuel 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, 505-665-8226, 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, 505-665-8226, thorp@lanl.gov Sue Kelly, SNL, 505-845-9770, smkelly@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov Ron Oldfield, SNL, 505-284-9153, raoldfi@sandia.gov
1.5.4.6	Post-Processing Environments	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov David Rogers, SNL, 505-844-5323, dhroger@sandia.gov
1.5.4.7	Common Computing Environment	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-8226, 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 Jim Mahoney, KCP, 816-997-5158, jmahoney@kcp.com

WBS	Title	Contact
1.5.5.2	User Support Services	Kim Cupps, LLNL, 925-423-7262, cupps2@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	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Kenneth Alvin, SNL, 505-844-9329, kfalvin@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 Mark Stalzer, 626-395-2521, stalzer@caltech.edu Susan Powell, 626-395-2909, spowell@cacr.caltech.edu
1.5.1.4-TRI-002	Purdue University, Center for Prediction of Reliability, Integrity and Survivability of Microsystems	Jayathi Murthy, 765-494-5701, jmurthy@ecn.purdue.edu Dawn Weisman, 765-494-8336, dweisman@purdue.edu
1.5.1.4-TRI-003	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, moin@stanford.edu Gianluca Iaccarino, 650-723-9599, iops@stanford.edu Deb Michael, 650-725-2077, debmich@stanford.edu
1.5.1.4-TRI-004	University of Michigan, Center for Radiative Shock Hydrodynamics	Paul Drake, 734-763-4072, rpdrake@umich.edu James Holloway, 734-936-3126, hagar@umich.edu
1.5.1.4-TRI-005	University of Texas, at Austin, Center for Predictive Engineering and Computational Sciences	Bob Moser, 512-471-0093, rmoser@mail.utexas.edu Chris Simmons, 512-232-2881, csim@colab.ices.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 the PSAAP, as listed below.

Academic Strategic Alliance Program:

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

The abovementioned universities will be closing out their Centers' activities funded under NNSA by end of FY10. They will continue to work on tasks that were in the FY09 IP but because a portion of the FY09 funding was cut, there are outstanding activities that will be completed in FY10.

Predictive Science Academic Alliance Program:

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

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 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 10^{11}/s, 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 FY10:

- 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 FY10:

- Determination of UQ analysis, in terms of confidence factors, on the extent to which the ballistic performance of the plates can be predicted by 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 FY11:

- 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.
- Deploy improved concentration-of-measure uncertainty bounds that exploit the hierarchical and multiscale structure of hypervelocity.
- Deploy a robust UQ-pipeline on a variety of NNSA platforms that supports the VTF and Eulerian simulators.
- Implement 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 civilian and defense systems. PRISM aims to significantly improve understanding of the long-term reliability of MEMS and 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.

Planned activities in FY10–FY 11

- 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 with enhanced potentials.
- Simulate integrated fluid-structure-electrostatic-thermal response of MEMS-related structures.
- Simulate dielectric charging of metal-insulator-metal (MIM) capacitor for single and multiple voltage cycles.
- Complete testing and optimization of MEMOSA FVM. Parallelize integrated MEMOSA FVM and material point method solvers.
- Develop first-principles based force-fields for molecular dynamics simulations, and use these force fields to study metal-dielectric contact.
- Perform V&V of FY10–11 target simulations, models, and software related to electrostatics, thermals and dielectric charging.
- Complete fluid-structure and fluid-structure-electrostatics UQ experiments to collect probability density function data for validation. Complete dielectric charging and electro-thermal experiments to provide validation data. Complete microsctructure characterization for PRISM device.
- Extend experimental micro-metrology framework for improved characterization of the geometry and material properties of packaged MEMS structures.

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 characterize the operability limits of an air-breathing hypersonic vehicle and associated propulsion system using predictive multi-physics simulations. The primary focus will be on the "unstart" failure mode triggered by thermal choking. 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. Current state-of-the-art simulation capabilities cannot predict these systems reliably, particularly near their operability limits.

World-class experimental facilities in Stanford's High Temperature Gasdynamics Laboratory (HTGL) will be used to conduct tightly integrated validation experiments for the key component physics and models. The development and implementation of UQ methods for very large systems will be an integral part of our effort at both the component and the system levels. Novel verification methods for high-fidelity simulations will also be developed and implemented. Stanford will leverage advanced computer science methods developed at Stanford to directly impact simulation tools and ensure scalability, program correctness, and portability to future platforms with very large numbers of cores.

Planned activities in FY10:

- Extend QMU analysis to include more detailed physical modeling and 2D and 3D simulations.
- Validate system-level simulations of Hyshot II scramjet configuration will be carried out using recent experiments at the High Enthalpy Shock Tunnel Goettingen, HEG, of the German Aerospace Center, DLR. Apply UQ analysis to quantify major uncertainties on heat release model.
- Implement combustion model into Charles LES solver, Extend combustion model to progress variable approach and enthalpy formulation, Verify and validate combustion model for both Joe and Charles solvers, including canonical problems and MMS for verification of reacting flow codes.
- Focus experimental program at HTGL on characterizing major uncertainties in H2/O2 system, including Pressure dependence of ignition delay times, Effect of impurities (e.g. O3, H-atoms), Known uncertainties in elementary reaction rates, and Behavior of high concentration mixtures.
- Apply both Charles (LES) and Joe (RANS) to the simulation of inlet/isolator experiment of Wagner/Clemens/Dolling, Stanford's gaseous jet in cross-flow (JIC) experiments, and the flame holding experiment at UTRC. A major focus will be characterizing the model-form uncertainties in RANS.
- Focus Stanford JIC experiments on characterizing mixing without reaction. Stanford experiments in the plasma physics laboratory will focus on characterizing the unstart bound due to blockage effects from jets.
- Use a Bayesian stochastic inversion to estimate the flight trajectory and speed (Hyshot II Flight data).
- Implement minimal Joe (non-reacting compressible flow solver with ideal gas) in the Liszt DSL.

Expected deliverables in FY10:

- End-to-end HyShot II simulation with combustion model, V&V'd combustion model, manufactured solution for hypersonic code validation, quantified impact of uncertainties on heat release.
- Initial results from Stanford PSAAP experiments, including: 1) the JIC experiments with mixing; 2) HTGL experiments on the H2/O2 system; 3) unstart experiments in the PPL; 4) UQ experiment in shockwave/boundary layer interaction.
- Analysis and timing comparisons for key kernels of the Joe code in DSL with comparison to production version.

Preliminary planned activities in FY11

- Focus on end-to-end HyShot II simulation with combustion, including quantified uncertainties in turbulence modeling.
- Develop support for the "weakly-intrusive" UQ methodology in the Joe solver.
- Continue coupled high-fidelity (LES, Charles) and lower-fidelity (RANS, Joe) simulation activity at the sub-system level to characterize the model-form uncertainties in RANS for hypersonic systems.
- Further DSL development and integration with Sequoia backend.

University of Michigan

The Center for Radiative Shock Hydrodynamics

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

- Developing software for radiation hydrodynamics to serve as a test bed for development and V&V of radiation hydrodynamics modeling elements.
- Developing a system for hierarchically validating the software.
- Extending an existing experimental effort, centered on radiative shocks, to obtain data and quantify uncertainties in the experiments.
- Simulating these experiments and assessing the predictive capability of the simulations.
- Establishing a doctoral program track for predictive science and engineering.

Planned activities in FY10:

- Simulate the signature experiment in 3D with version 2 of the CRASH code.
- Perform first full cycle of assessment of predictive capability based on 1D simulations. This will involve screening for variable selection, other forms of dimension reduction, simulations over distributions of values including variations reflecting uncertainty in constants of nature, a Bayesian analysis involving a comparison of simulation results to distributions of actual data, and other steps.
- Perform experiments to quantify the behavior of the Be plasma that drives the radiative shock.
- Modify the BATSRUS/CRASH code to: i) include electron heat conduction and separate electron and ion temperatures, ii) function with radiation and AMR in 1D, 2D, and 3D, iii) include multigroup radiation diffusion, iv) couple with PDT for multigroup SN radiation transport, and v) determine the multigroup opacities and other parameters needed for the previous items.
- Modify the PDT radiation transport code to: i) include a 2D axisymmetric (r,z) capability, ii) improve the interface with BATSRUS, iii) include a diffusion preconditioner within its Krylov iterative method, iv) take advantage of improvements in the underlying library as they are developed by the STAPL team.
- Work actively on the underlying physics of the radiative shocks of interest, to better develop the expert judgment that is essential for assessing predictive capability.

- Continue to develop, implement, and document numerous test problems ranging from unit tests to multi-physics tests.

Expected deliverables in FY10:

- A report of a 3D simulation of the signature experiment with version 2 of the CRASH code
- A software report, as part of the annual technical report, detailing the accomplishments in code development.
- A UQ report, as part of the annual technical report, describing the results of the first cycle of assessment of predictive capability.
- An experiments report, as part of the annual technical report, describing the analysis of the experiments to assess variability and the experiment design for the Be plasma experiment.
- A theoretical description of the system of interest, as part of the annual technical report, describing progress in fundamental understanding and semi-analytic modeling of this system.

Preliminary planned activities in FY11:

- Improve the capabilities of the simulation, with priorities set by what is learned in the assessment of predictive capability. This may involve routine use of steady-state adjoint methods for grid adaptation, implementation of separate grid refinement for hydrodynamics and radiation, improvements to interface tracking, or other developments.
- Identify, plan, and perform three experiments based on areas of maximum importance determined by the assessment of predictive capability.
- Continue to improve the fundamental understanding of the physical system of interest.
- Perform an assessment of predictive capability involving 2D simulations.

University of Texas

The Center for Predictive Engineering and Computational Sciences

The goal of the PECOS Center is to develop next generation 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 via pyrolysis, 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 FY10:

- Calibrate and validate radiation, chemistry and data reduction models via shock tube data
- Develop and implement large-scale variants of Stochastic Newton method
- Develop new, adjoint-enable hypersonic flow code
- Calibrate and validate ablation models using new and legacy data
- Calibrate and validate turbulence models using legacy and DNS data, including model uncertainty models

Expected deliverables in FY10:

- Implementation of initial adjoint-enabled ablation model and hypersonic flow codes
- Calibration and validation of component level models in all modeling domains
- Completion of coupled system simulations of the reentry capsule with uncertainty propagation including first quantification of uncertainty in heat fluxes and ablation rates using available estimates of input uncertainties

Preliminary planned activities in FY11:

- Continue working on the model development, implementation and code verification
- Continue working on model validation using updated data across all domains
- Develop and implement intrusive reduction methods into sampling for Bayesian inversion
- Develop and implement intrusive reduction methods into uncertainty propagation software
- Develop adjoint-enable modeling codes across all domains
- Conduct coupled system simulations of the reentry capsule using uncertainty propagation including quantification of uncertainty in heat fluxes and ablation rates using calibrated input uncertainties

Appendix E.
ASC Obligation/Cost Plan

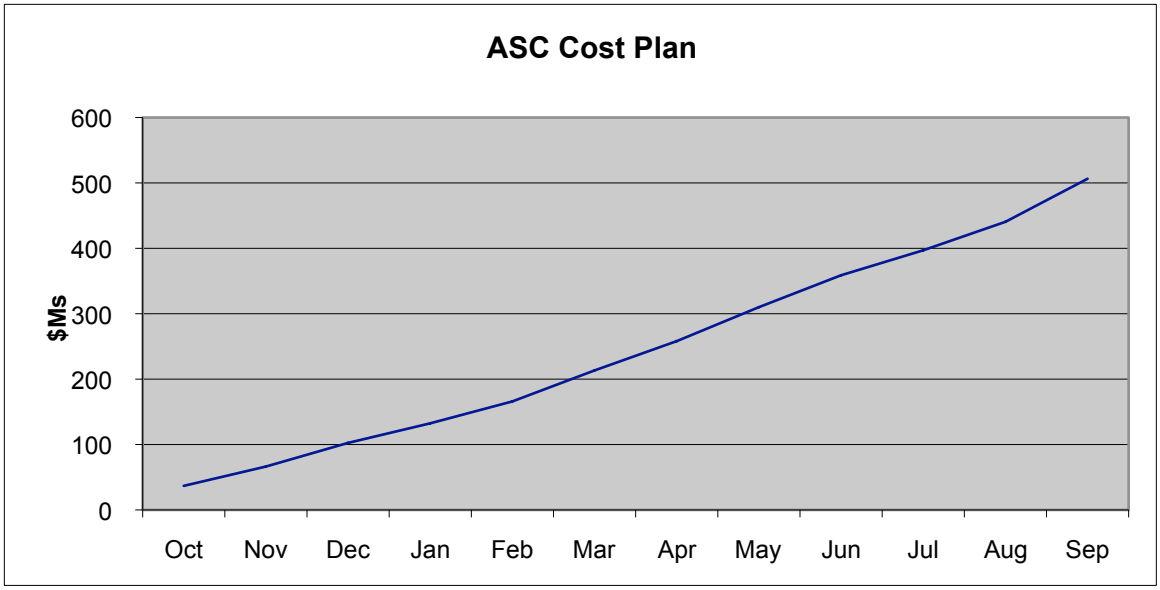

Figure D-1. ASC obligation/cost plan for FY10.

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