LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-676591

Advanced Simulation and Computing Fiscal Year 2016 Implementation Plan, Version 0

M. McCoy, B. Archer, B. Hendrickson

August 27, 2015

Disclaimer

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

National Nuclear Security Administration

Advanced Simulation and Computing
FY16 IMPLEMENTATION PLAN

Version 0

August 27, 2015

NA-ASC-127R-15-Rev.0-IP

LLNL-TR-676591

National Nuclear Security Administration

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	*Michel McCoy, Bill Archer, Bruce A. Hendrickson*	*initial*	*Doug Wade, Thuc Hoang*	*8/27/15*	*Initial release*
1					*Update after signed budget*

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

National Nuclear Security Administration

Contents

I. Overview

The Stockpile Stewardship Program (SSP) is an integrated technical program for maintaining the safety, surety, and reliability of the U.S. nuclear stockpile. The SSP uses nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program requires the continued use of experimental facilities and programs, and the computational capabilities to support these programs.

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

As the program approaches the end of its second decade, ASC is focused on increasing 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 (sufficient resolution, dimensionality, and scientific details), and quantifying critical margins and uncertainties. Resolving each issue requires increasingly difficult analyses because the aging process has progressively moved the stockpile further away from the original test base. Where possible, the program also enables the use of high performance computing (HPC) and simulation tools to address broader national security needs, such as foreign nuclear weapon assessments and counternuclear terrorism.

In Prague (2009), and more recently Berlin (2013), President Obama articulated his vision of a world without nuclear weapons. These reductions will be made, however, while ensuring that the U.S. maintains a safe, secure, and effective strategic deterrent for as long as such weapons exist.

The 2010 *Nuclear Posture Review Report* (NPR) codified the role of the National Nuclear Security Administration (NNSA) in maintaining the deterrent. In areas essential for stockpile life extensions and stewardship, key investments have been made to:

- Strengthen the science, technology, and engineering base needed for conducting weapon system LEPs
- Mature advanced technologies to increase weapons surety
- Qualify weapon components and certify weapons without nuclear testing

- Provide annual stockpile assessments through weapons surveillance

This strategy includes developing and sustaining high-quality scientific staff, as well as supporting computational and experimental capabilities. These components constitute the foundation of the nuclear weapons program.[1]

The continued success of the SSP and LEPs is predicated upon the ability to credibly certify the stockpile, without a return to underground nuclear tests (UGTs). Shortly after the nuclear test moratorium entered into force in 1992, the Accelerated Strategic Computing Initiative (ASCI) was established to provide the underpinning simulation capability to support stockpile certification. While computing and simulation have always been essential to the success of the nuclear weapons program, the program goal of ASCI was to execute NNSA's vision of using these tools in support of the stockpile mission. The ASCI Program was essential to the success of the SSP, providing critical nuclear weapons simulation and modeling capabilities. Now designated as the ASC Program, the mission remains the same: provide the simulation and computational capabilities that underpin the ability to maintain a safe, secure, effective nuclear weapon stockpile, without a return to UGTs. Where possible, the program also enables the use of these tools and operational infrastructure to address broader national security needs, such as foreign nuclear weapon assessments and counternuclear terrorism.

The capabilities that the ASC Program provides at the national laboratories play a vital role in the Nuclear Security Enterprise, and are necessary for fulfilling the stockpile stewardship and life extension requirements outlined for NNSA in the NPR report. The Program develops modern simulation tools that provide insights into stockpile aging issues, provide the computational and simulation tools that enable designers and analysts to certify the current stockpile and life-extended nuclear weapons, and inform the decision making process when any modifications in nuclear warheads or the associated manufacturing processes are deemed necessary. Furthermore, ASC is enhancing the predictive simulation capabilities that are essential to evaluate weapons effects, design experiments, and ensure test readiness.

The ASC Program continues to improve its unique tools to solve stockpile problems—with a focus on sufficient resolution, dimensionality, and scientific detail—to enable Quantification of Margins and Uncertainties (QMU) and to resolve the increasingly difficult analyses needed for stockpile stewardship. The needs of the Directed Stockpile Work (DSW) and major modernization programs also drive the requirements for simulation and computational resources. These requirements include planned LEPs, stockpile support activities, and mitigation efforts against the potential for technical surprise. All of the weapons within the current stockpile are in some stage of the life-extension process. The simulation and computational capabilities are crucial for successful execution of these lifetime extensions and for ensuring NNSA can certify these life-extended weapons without conducting a UGT.

Specific work activities and scope contained in this Implementation Plan (IP) represent the full-year annual operating plan for FY16. The Initial IP effective October 1, 2015,

[1] *2010 Nuclear Posture Review Report*, April 2010, p. 42.

should be consistent with: 1) the Department's Base Table when operating under a Continuing Resolution (CR); and 2) the final signed appropriation or full-year CR once enacted.

II. Corporate Program Goals

Preliminary targets are subject to change based on a final, enacted budget.

Program or Project Name	Performance Measure/ Indicator Title and Description	FY2016 Target	Endpoint Target
Advanced Simulation and Computing Program	Reduced Reliance on Calibration	53%	100% (FY2024)

The contractor's *Performance Evaluation Plan* contains multisite targets that can be identified by the associate deputy administrator as base or stretch goals.

There are no multisite targets (MST) for ASC.

Along with the Contributing Factors and Site Specific Outcomes outlined in the *Performance Evaluation Plan*, the contractor's performance will be evaluated against the NNSA's *Strategic Plan*, NNSA performance priorities and deliverables, program execution plans, work authorizations (WAs), and other key inputs (for example, multiyear strategic objectives). In evaluating overall performance on the FY16 milestones, the contractor shall receive adjectival ratings "Excellent," "Very Good," "Good," "Satisfactory," or "Unsatisfactory" based on Federal Acquisition Regulation Subpart 16.401(e)(3).

At a minimum, all management and operating (M&O) sites are expected to perform at the satisfactory level documented in the Strategic *Performance Evaluation Plan* for each site. If not stated specifically in the Strategic *Performance Evaluation Plan*, satisfactory performance includes achieving all milestones and/or keeping NNSA informed of obstacles to achieving milestones that may arise due to the scientific discovery nature of the ASC work; meeting all reporting requirements; engaging in productive and constructive collaboration with other ASC partner sites especially to achieve joint milestones and to achieve joint, collaborative, scientific goals; productive and constructive peer review of ASC partners; constructive participation in ASC meetings and reviews; professional interactions especially between management and NNSA; and cost-effective management of ASC funds and facilities.

III. Major Activities

The statutory objective of the SSP is to ensure a high level of confidence in the safety, reliability, and performance of weapons in the nuclear stockpile. The ASC Program provides high-end simulation capabilities to meet the requirements of the SSP, and it includes weapon codes, computing platforms, and supporting infrastructure. The ability to model the extraordinary complexity of nuclear weapons systems is essential to maintaining confidence in the performance of the aging stockpile without underground testing. The ASC Program underpins the Annual Assessment Review (AAR) of the stockpile and is an integrating element of the Predictive Capability Framework (PCF), as described in the FY16 Stockpile Stewardship and Management Plan[2]. ASC also provides critical capabilities informing efforts to extend the life of the nuclear stockpile.

The ASC capabilities are also used to address areas of national security in addition to the U.S. nuclear stockpile. Through coordination with other government agencies and other organizations within NNSA, ASC plays important roles in supporting nonproliferation, emergency response, nuclear forensics, and attribution activities.

The ASC Program is split into six subprograms:

Integrated Codes (IC) subprogram contains the mathematical descriptions of the physical processes of nuclear weapon systems and function. Combined with weapon-specific input data created by the nuclear weapons designers and engineers, this allows detailed simulations of nuclear weapons performance assessment without the need for underground nuclear testing. The IC subprogram funds the critical skills needed to develop, maintain, and advance the capabilities of the large-scale integrated simulation codes needed for the following SSP and DSW activities: annual assessment; LEP design, qualification, and certification; SFI resolution; and safety assessments to support transportation and dismantlement. In addition, these capabilities are necessary for a host of related requirements such as nuclear counter-terrorism efforts (for example, nuclear forensics, foreign assessments, and device disablement techniques).

The **Physics and Engineering Models** (PEM) subprogram provides the models and databases used in simulations supporting the U.S. stockpile. These models and databases describe a great variety of physical and engineering processes occurring in a nuclear weapon over its full lifecycle. The capability to accurately simulate these processes is required for annual assessment; design, qualification, and certification of warheads undergoing LEPs; resolution (and in some cases generation) of SFIs; and the development of future stockpile technologies. The PEM subprogram is closely linked to the Science program, which provides the experimental data that informs development of new models used in simulation codes.

[2] U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2016 Stockpile Stewardship and Management Plan*, Report to Congress, March 2015.

The **Verification and Validation** (V&V) subprogram provides evidence that the models in the codes produce mathematically correct answers that reflect physical reality. The V&V subprogram funds the critical skills needed to apply systematic measurement, documentation, and demonstration of the ability of the models and codes to predict physical behavior. The V&V subprogram is developing and implementing uncertainty quantification (UQ) methodologies as part of the foundation for the QMU process of weapons assessment and certification. The V&V subprogram also drives software engineering practices to improve the quality, robustness, reliability, and maintainability of the codes that evaluate and address the unique complexities of the stockpile. As nuclear test data is becoming less relevant with an aging stockpile and as weapons designers with test experience leave the nuclear security enterprise, it has become increasingly important that the codes are verified and validated so that future generations of designers are confident in the use of these foundational tools.

V&V efforts and predictive capability assessments will continue to increase the ASC Program's ability to address complex safety and engineering issues within the nuclear weapons stockpile. With major modifications to adapt existing codes to future hardware (a major focus of the IC subprogram) and development of new codes (a primary focus of the Advanced Technology Development and Mitigation (ATDM) subprogram), V&V will ensure the modifications and new codes are subjected to thorough V&V activities. This will be a major focus area for the V&V subprogram.

The **Advanced Technology Development and Mitigation** subprogram includes laboratory code and computer engineering and science projects that pursue long-term simulation and computing goals relevant to both exascale computing and the broad national security missions of the NNSA.

ASC capabilities that support the DSW mission are beginning to stall, as HPC technologies are evolving to radically different and more complex (many-core, heterogeneous) architectures. Efficiency of the integrated design codes (IDC) is falling significantly when they are used on the latest HPC platforms, and this trend is expected to accelerate and spread unless mitigated. Three major challenges to address through investment in this subprogram include: 1) the radical shift in computer architecture, 2) maintenance of the current millions of lines of IDC that took more than a decade to develop and validate, and 3) sustainment and adaptation of current capabilities as evolving computer technologies become increasingly disruptive to the broad national security missions of NNSA.

There are two focus areas for investment. The Next Generation Code Development and Application project is focused on long-term research that investigates how future code development must address new HPC challenges of massive, heterogeneous parallelism using new programming models and data management techniques developed through co-design of applications and systems. The Next Generation Architecture and Software Development project is focused on long-term computing technology research of extreme, heterogeneous architectures to mitigate the impact of architectural change and advance its capabilities for ASC simulation codes.

The ATDM subprogram tackles the most critical subset of issues occurring during this period of disruptive change in HPC architectures to continue the current level of support to the DSW mission.

The **Computation Systems and Software Environment (CSSE)** subprogram builds the computing systems needed for weapons simulations. Since requirements of the ASC codes drives the need to achieve its predictive capability goals, the ASC Program must continue to invest in and consequently influence the evolution of computational environments. Along with the powerful commodity technology (CT) and advanced technology (AT) systems that the program fields, the supporting software infrastructure deployed on these platforms includes many critical components, from system software to input/output (I/O), storage and networking, and post-processing visualization and data analysis tools.

The **Facility Operations and User Support (FOUS)** subprogram provides the facilities and services required to run nuclear weapons simulations. Facility operations includes physical space, power, and other utility infrastructure, and local area/wide area networking for local and remote access, as well as system administration, cyber-security, and operations services for ongoing support. User support includes computer center hotline and help-desk services, account management, Web-based system documentation, system status information tools, user training, trouble-ticketing systems, common computing environment (CCE), and application analyst support.

These six subprograms (IC, PEM, V&V, ATDM, CSSE, and FOUS) all contribute to a cohesive set of program deliverables. Highlights of the FY16 major activities for the ASC Program include:

- Complete work on defining early initial conditions for boost; begin updating the IDC with results.
- Deploy CT systems and complete the Trinity system for the tri-labs' production computing environment to address stockpile stewardship issues and to advance predictive science.
- Continue the development of the ATDM subprogram, to mitigate the impact of new computer architectures on current code capabilities.
- Expand the predictive capability assessment suites to include additional UGTs, hydrodynamic tests, and scaled experiments
- Maintain full baselines for all stockpile systems and use these baselines to improve the fidelity of annual stockpile assessments
- Coordinate and collaborate HPC technology research, development, and engineering activities in partnership with the Department of Energy (DOE)/Advanced Scientific Computing Research (ASCR) office, to advance technologies that will eventually enable procurement of an exascale-class HPC platform.

The ASC budgetary increase in FY16 provides the funding required to facilitate transitioning the integrated codes to work efficiently on emerging high-performance

computers, develop next-generation codes, and maintain computing resources and facilities. These capabilities are necessary to inform the annual assessment of the nuclear stockpile. The drivers of the ASC Program that require these budgets include the following:

- The Nuclear Weapons Council approved the Long Range Stockpile Sustainment Strategic Plan, a key aspect of which is the "3+2 Strategy." Supporting the 3+2 strategy requires further developed simulation and computing capabilities to enable progress in understanding energy balance, boost, and improved equations of state for materials of interest.
- Annual assessments, LEPs, and SFIs require responsive modeling and simulation capabilities to better understand the impact of environmental and system conditions, including aging and the resolution of historical nuclear test anomalies.
- Investing in physics improvements in the IDC will open design options for subsystem components for future LEPs.

The ASC computing capabilities are the key integrating mechanisms across the nuclear weapons program through the IDC. The assessment of the nation's stockpile requires high-fidelity physical models. The IDC support design studies, maintenance analyses, the annual assessment reports, LEPs, SFIs, and weapons dismantlement activities. The IDC contain the mathematical descriptions of the physical processes of nuclear weapon systems and function. Combined with weapon-specific input data created by the nuclear weapons designers and engineers, the IDC allow detailed simulations of nuclear weapons performance assessment, without the need for underground nuclear testing. Since the 1992 nuclear weapons testing moratorium, the IDC embody the repository of data from experiments conducted at the NNSA's high energy density facilities and legacy UGTs, as well as the accumulated experience of the DSW user community. The IDC currently perform well for general mission-related activities; however, as the stockpile is life extended and aging takes the current stockpile further away from the data collected from UGTs, maintaining the nuclear weapons stockpile will require IDC that enhance prediction and use HPC resources more effectively.

A strategic driver for simulation and computing investment is the global shift in fundamental computing architecture. ASC capabilities that support the DSW mission are beginning to experience the effects of obsolescence as HPC technologies continue to advance and evolve to radically different and more complex (with massively concurrent cores, heterogeneous, and memory limiting) architectures. Maintaining currency with the commercial information technology sector will advance high-fidelity physics modeling capabilities required to maintain a credible deterrent and will address additional mission needs in non-proliferation, emergency response, and nuclear forensics and attribution programs. To address this strategic driver, ASC is redirecting resources to minimize the disruptive mission impact of this change in HPC.

The ASC Program has developed a strategy for acquiring the advanced computing technologies needed to support current and future stockpile work that fully recognizes the need for the acquisition of exascale computing capabilities in the future. The ASC

Program's approach to advancing HPC technologies in this strategy is scoped to contribute to the foundation for an exascale supercomputer capability for the nation. The new ATDM subprogram consolidates the investments Congress directed in FY14 for exascale into a unified effort to tackle challenges facing ASC in its support to stockpile stewardship and upon which future efforts can build. Since the technical problems facing the program today are similar issues but at lower scale—issues that exascale will need to overcome to be successful—investments in ATDM advance both exascale technologies and stockpile computing effectiveness.

IV. Funding Guidance

To support the scope of work contained in this IP, funding will be distributed through the existing Approved Funding Program (AFP) process. The AFP is adjusted on an as needed basis for the execution of Congressionally approved programs, projects, or activities. Specific work activities are authorized via this document, with incremental funding changes made through the AFP, and authorized via WAs and obligated via formal contract modification.

Program/Operational Control Level	President's FY2016 Budget Request	FY2016 CR Operating Target	FY2016 Enacted/Full-Year CR	Difference between Request versus Enacted
ASC	$623M			

V. Description of Planned Activities

The purpose of this IP is to outline key work requirements to be performed and to control individual work activities within the scope of work. Contractors may not deviate from this plan without a revised WA or subsequent IP.

Specific quantifiable subprogram deliverables are negotiated and/or updated during an annual process to document and track subprogram Level 2 Milestones. Successful progress toward completing these milestones is tracked on a quarterly basis. Progress towards completion of subprogram deliverables contribute toward an aggregate assessment of the Campaign's progress toward a quantifiable total number of deliverables for the current fiscal year.

Annual performance expectations for each M&O contractor outlined in this document will be considered in determining the contractor's performance rating and fee earned through the NNSA Corporate Performance Evaluation Process (CPEP).

The table below lists the current ASC Level 2 Milestones for FY16. A more comprehensive list (including milestone description and grading/exit criteria) is included in the individual subprogram detail in the appendices.

Table V-1. ASC Level 2 Milestone for FY16[3]

Sub-Program	ID#	Milestone Title	Complete Date	Site
IC	TBD	Evaluate the Use of Arbitrary Lagrangian-Eulerian/Adaptive Mesh Refinement in Full System Modeling, Part 1	12/31/15	LLNL
IC	TBD	Demonstrate a New Embedded Smooth Particle Hydrodynamic Capability	6/30/16	LLNL
IC	TBD	Deliver within the Nuclear Performance Code System, an Initial Capability for the Simulation of Output	9/30/16	LLNL
IC	TBD	Demonstration of Eulerian Application Project Structural Upgrade for Future Architectures	9/30/16	LANL

[3] Factors such as FY16 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

Sub-Program	ID#	Milestone Title	Complete Date	Site
IC	TBD	Demonstration by the Lagrangian Application and Setup Project of a 3D Simulation Capability	9/30/16	LANL
IC, PEM	TBD	Delivery of III-IV XYCE (RAMSES) Compact Models	9/30/16	SNL
IC, PEM	TBD	Create a Two-Way Coupled Aero-Structural Capability for Re-Entry Environments	9/30/16	SNL
PEM	TBD	Multiphase Plutonium Kinetics Framework	9/3016	LLNL
PEM	TBD	New Plutonium Equation of State Tables Based on Recent Data and Simulations	9/30/16	LLNL
PEM	TBD	Dynamic Fracture Model for Brittle Metal Phases	9/30/16	LANL
PEM	TBD	Integrated Material Models for Equation of State and Transport Properties for Aluminum Supporting Uncertainty Quantification	9/30/16	SNL
PEM	TBD	Coupled Computational Fluid Dynamics/Ablation Simulation of Hypersonic Vehicles	9/30/16	SNL
PEM, V&V	TBD	Develop a Multiphysics Capability for Predicting the Evolving Material State through the Manufacturing Process of a Gas Transfer System Reservoir	9/30/16	SNL
V&V	TBD	Assess Uncertainty of Secondary Performance Using a Proposed Modern Methodology for Uncertainty Quantification	9/30/16	LLNL
V&V	TBD	Diagnostic Data Analysis, Uncertainty, and Modeling in an ASC Era	9/30/16	LANL
V&V	TBD	Provide Enhanced Feature Coverage for Assessment of Application-Specific Model Credibility along with Suitable Tool Verification and Testing	6/30/16	SNL

Sub-Program	ID#	Milestone Title	Complete Date	Site
V&V	TBD	Blind Validation of Stochastic Effects in Directed Stockpile Work III-IV Heterojunction Bipolar Transistors	9/30/16	SNL
ATDM	TBD	Demonstration of at Least One Modular Hydrodynamic Capability within the Computer Science Toolkit*	9/30/16	LLNL
ATDM	TBD	Deliver Strategic Plan and Initial Scalability Assessment by Advanced Architecture and Portability Specialists Team	9/30/16	LLNL
ATDM, IC, CSSE	TBD	Next-Generation Code: Plan for Rapid Development	9/30/16	LANL
ATDM	TBD	Demonstrate Performance Portability and Embedded Analysis Infrastructure	9/30/16	SNL
ATDM, IC, CSSE	TBD	Demonstrate Integration of Performance Improvements Identified via Proxy Applications into Application Codes	9/30/16	LLNL, LANL, SNL
CSSE	TBD	Modernization and Expansion of LLNL Archive Disk Cache	3/31/16	LLNL
CSSE	TBD	Deploy Commodity Technology System-1 Tripod Operating System Software Test Bed Scalable Unit (Hype)	6/30/16	LLNL
CSSE, FOUS	TBD	Trinity-Haswell High Performance Computing System Delivery	12/31/15	LANL
CSSE	TBD	Demonstrate and Evaluate Advanced Analysis, Visualization, and Input/Output Capabilities for the SIERRA Toolkit	6/30/16	SNL
CSSE	TBD	Evaluate Impact of Advanced Memory Architectures on ASC Codes	9/30/16	SNL
FOUS	TBD	Further Automate Planned Cluster Maintenance to Minimize System Downtime during Maintenance Windows	9/30/16	LLNL
FOUS	TBD	Building 654 Beneficial Occupancy	9/30/16	LLNL
FOUS	TBD	Deploy a Common Computing Environment for Commodity Technology System-1 Platforms	9/30/16	LLNL, LANL, SNL

Sub-Program	ID#	Milestone Title	Complete Date	Site

* The long description for this Level 2 milestone is available upon request from the ASC Program Office.

VI. Reporting Requirements

The following systems and processes for program management and control of the ASC Program are in place:

- **Quarterly Program Reviews**. M&O contractors report Level 2 milestone status to Headquarters (HQ) using the Milestone Reporting Tool. In preparation for each quarterly review, each site and Federal Program Manager will assess the status of each milestone by providing a score (Green, Yellow, Red, Blue, or Black). In addition, supporting details for the assessment of each milestone will be reported and any programmatic risk will be identified.
- **Monthly Financial Reporting**. Monthly cost/financial expenditure data will be reported by the Office of Planning, Programming, Budgeting and Evaluation, NA-MB, and available to the NNSA Program Managers.
- **Other NNSA Program Reviews**. Special technical and program reviews requested by NNSA Program Managers and other senior NNSA officials for oversight and program management responsibilities will be supported by the sites.
- **Bi-Weekly Subprogram Teleconference**. Federal Program Managers conduct bi-weekly teleconferences to discuss upcoming meetings and to provide an opportunity to exchange information of programmatic and technical interest and need.
- **Program Change Control**. Change control for program activities conducted within this IP will be managed and tracked on the Revision Summary at the front of this document.
- **Budget Control Levels**. The budget control level allows the federal Program Manager to shift funding within the ASC Campaign subprograms. All requests to shift resources between subprograms must be approved in advance by the appropriate NNSA Program Manager.
- **Corporate Performance Evaluation Process (CPEP)**. Each NNSA M&O contractor is evaluated utilizing the individual contract's Strategic Performance Evaluation Plan. Program Managers are required to establish the expectations for the M&O contractor(s) via this IP and associated WAs. The annual evaluation of each M&O contractor is performed per the CPEP Process Policy Guide. The Federal Program Managers provided quarterly evaluations, which are included in the annual Performance Evaluation Report produced by the NNSA Field Office.

VII. Key Execution Year Reference Documents

The following documents are incorporated by reference:

- *ASC Computing Strategy* (May 2013)
- *ASC Right Size* (October 2010)
- *ASC Business Plan* (July 2015)
- MOU between DOE Office of Science (SC) and DOE NNSA Office of Defense Programs for the coordination of exascale activities (April 2011)
- *ASC Program Management Plan* (July 2010)
- AFP Input sheet and regular monthly financial plan adjustments, including Work Breakdown Structure (WBS), Budget and Reporting (B&R) code, and other necessary information for each site in the monthly AFP updates

VIII. Major Risks and Handling Strategies

A number of factors must operate in concert to ensure the work proceeds as planned. Deviation from any one of these factors may cause delays in milestone schedules, reductions in scope, or increased technical risks and uncertainties. Technical risks specific to an individual milestone are covered in the individual subprogram appendices to this document.

IX. Points of Contact

Headquarters (HQ) Programmatic Points of Contact

Doug Wade
NA-114, Acting Director
Tele.: 202-587-7507
FAX: 202-586-7754
douglas.wade@nnsa.doe.gov

HQ IP Focal Point

Thuc Hoang
NA-114
Tele.: 202-586-7050
FAX: 202-586-7754
thuc.hoang@nnsa.doe.gov

Laboratory ASC Executives

Michel McCoy, LLNL, 925-422-4021, mccoy2@llnl.gov

Bill Archer, LANL, 505-665-7235, barcher@lanl.gov

Bruce Hendrickson, SNL, 505-845-7599, bahendr@sandia.gov

Laboratory Programmatic Points of Contact

WBS	Title	Contact
1.2.3.1	Integrated Codes	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Jerry Brock, LANL, 505-665-3210, jsbrock@lanl.gov Scott Hutchinson, SNL, 505-845-7996, sahutch@sandia.gov
1.2.3.1.1	Engineering and Physics Integrated Codes	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Jerry Brock, LANL, 505-665-3210, jsbrock@lanl.gov Scott Hutchinson, SNL, 505-845-7996, sahutch@sandia.gov
1.2.3.1.2	Specialized Codes and Libraries	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Jerry Brock, LANL, 505-665-3210, jsbrock@lanl.gov Scott Hutchinson, SNL, 505-845-7996, sahutch@sandia.gov
1.2.3.1.3	Applications and Algorithms Research	Chris Clouse, LLNL, 925-422-4576, clouse@llnl.gov Jerry Brock, LANL, 505-665-3210, jsbrock@lanl.gov Scott Hutchinson, SNL, 505-845-7996, sahutch@sandia.gov
1.2.3.1.4	Applications Research for Next-Generation Platforms	Chris Clouse, LLNL, 925-422-4576, clouse@llnl.gov Jerry Brock, LANL, 505-665-3210, jsbrock@lanl.gov Scott Hutchinson, SNL, 505-845-7996, sahutch@sandia.gov
1.2.3.2	Physics and Engineering Models	Dennis McNabb, LLNL, 925-423-0749, mcnabb3@llnl.gov Mark Schraad, LANL, 505-665-3946, schraad@lanl.gov Jim Redmond, SNL, 505-844-3136, jmredmo@sandia.gov
1.2.3.2.1	Materials Response: Equation of State, High Explosives, and Mix and Burn	Dennis McNabb, LLNL, 925-423-0749, mcnabb3@llnl.gov Mark Schraad, LANL, 505-665-3946, schraad@lanl.gov Jim Redmond, SNL, 505-844-3136, jmredmo@sandia.gov
1.2.3.2.2	Transport, Plasmas, Atomic, Nuclear	Dennis McNabb, LLNL, 925-423-0749, mcnabb3@llnl.gov Mark Schraad, LANL, 505-665-3946, schraad@lanl.gov Jim Redmond, SNL, 505-844-3136, jmredmo@sandia.gov

WBS	Title	Contact
1.2.3.2.3	Engineering Science, Environments, and Response	Jim Redmond, SNL, 505-844-3136, jmredmo@sandia.gov
1.2.3.2.4	Integrated Modeling and Applications	Dennis McNabb, LLNL, 925-423-0749, mcnabb3@llnl.gov Mark Schraad, LANL, 505-665-3946, schraad@lanl.gov Jim Redmond, SNL, 505-844-3136, jmredmo@sandia.gov
1.2.3.3	Verification and Validation	Joe Sefcik, LLNL, 925-423-0671, sefcik1@llnl.gov Fred Wysocki, LANL, 505-667-6543, fjw@lanl.gov Mike Chiesa, SNL, 925-294-2103, chiesa@sandia.gov
1.2.3.3.1	Verification and Validation Methods	Joe Sefcik, LLNL, 925-423-0671, sefcik1@llnl.gov Fred Wysocki, LANL, 505-667-6543, fjw@lanl.gov Mike Chiesa, SNL, 925-294-2103, chiesa@sandia.gov
1.2.3.3.2	Verification and Validation Assessments	Joe Sefcik, LLNL, 925-423-0671, sefcik1@llnl.gov Fred Wysocki, LANL, 505-667-6543, fjw@lanl.gov Mike Chiesa, SNL, 925-294-2103, chiesa@sandia.gov
1.2.3.3.3	Data Validation, Archiving, Software Quality Assurance, and Training	Joe Sefcik, LLNL, 925-423-0671, sefcik1@llnl.gov Fred Wysocki, LANL, 505-667-6543, fjw@lanl.gov Mike Chiesa, SNL, 925-294-2103, chiesa@sandia.gov
1.2.3.4	Advanced Technology Development and Mitigation	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Mark Anderson, LANL, 505-667-4772, mcanderson@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov
1.2.3.4.1	Next-Generation Code Development and Application	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Mark Anderson, LANL, 505-667-4772, mcanderson@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov
1.2.3.4.2	Next-Generation Architecture and Software Development	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Mark Anderson, LANL, 505-667-4772, mcanderson@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov

WBS	Title	Contact
1.2.3.4.3	Future High Performance Computing Technologies	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Mark Anderson, LANL, 505-667-4772, mcanderson@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov
1.2.3.5	Computational Systems and Software Environment	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Mike Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.1	Commodity Technology Systems	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.2	Advanced Technology Systems	Bronis de Supinski, LLNL, 925-422-1062, desupinski1@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.3	System Software and Tools	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Mike Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.4	I/O, Storage Systems, and Networking	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Mike Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.5	Post-Processing Environments	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Mike Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.6	Next-Generation Computing Technologies	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Mike Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.6	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Paul Weber, LANL, 505-606-2125, pmw@lanl.gov Tom Klitsner, SNL, 505-844-1901, tklitsn@sandia.gov

WBS	Title	Contact
1.2.3.6.1	Collaborations	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Paul Weber, LANL, 505-606-2125, pmw@lanl.gov David Womble, SNL, 505-845-7471, dewombl@sandia.gov
1.2.3.6.2	System and Environment Administration and Operations	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Paul Weber, LANL, 505-606-2125, pmw@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.2.3.6.3	Common Computing Environment	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Paul Weber, LANL, 505-606-2125, pmw@lanl.gov Joel Stevenson, SNL, 505-845-1347, josteve@sandia.gov
1.2.3.6.4	Special Purpose Facilities, Systems, Operations, and Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Paul Weber, LANL, 505-606-2125, pmw@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
University of Utah, The Uncertainty Quantification-Predictive Multidisciplinary Simulation Center for High Efficiency Electric Power Generation with Carbon Capture		Prof. Philip Smith philip.smith@utah.edu (801) 585-3129 Jeri Schryver (Administrative Manager) 801-581-8712 jeri@eng.utah.edu
University of Illinois, Urbana-Champaign, Center for Exascale Simulation of Plasma-Coupled Combustion		Prof. William Gropp wgropp@illinois.edu (217) 244-6720 Richard D. Mihm (Manager of Accounting) 217-244-8886 rmihm@illinois.edu
Stanford University, Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment		Prof. Gianluca Iaccarino jops@stanford.edu (650) 906-0302 Joanne H. Lee (Research and Financial Administrator) 650-721-6283 Joanne.Lee@stanford.edu

WBS	Title	Contact
University of Florida, Center for Compressible Multiphase Turbulence		Prof. S. 'Bala' Balachandar bala1s@ufl.edu (352) 392-8909 Melanie DeProspero (Assistant to Research Group) 352-846-3018 Mel7703@ufl.edu
Texas A&M University, Center for Exascale Radiation Transport		Prof. Jim Morel morel@tamu.edu (979) 845-4161 Joan D. Birdwell (Senior Project Administrator II) 979-862-7984 jbirdwell@tamus.edu
University of Notre Dame, Center for Shock Wave Processing of Advanced Reactive Materials		Prof. Samuel Paolucci paolucci@nd.edu (574) 631-8110 John H. Engel (Senior Accountant) 574-631-7088 John.H.Engel.3@nd.edu

X. Approvals

The undersigned acknowledge that they have reviewed the ASC IP and agree with the information presented within this document. Changes to this IP will be coordinated with, and approved by, the undersigned, or their designated representatives.

Approved by:

Douglas Wade,
NNSA ASC Program Director, Acting 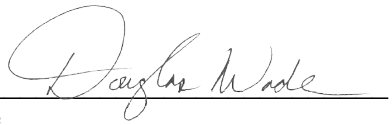 8/27/15
Signature Date

Concurrence:

Michel McCoy, LLNL ASC Executive	*Concurrence received*	8/19/15
Bill Archer, LANL ASC Executive	*Concurrence received*	8/25/15
Bruce A. Hendrickson, SNL ASC Executive	*Concurrence received*	8/20/15

Appendix A: Key Terms

The following definitions and explanations are for terms and acronyms relevant to the content presented within this document and its appendices.

3D	Three Dimensional
AAPS	Advanced Architecture and Portability Specialist
AAR	Annual Assessment Review
ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AFP	Approved Funding Program
ALE	Arbitrary Lagrangian-Eulerian
AMD	Advanced Micro Devices
AMR	Adaptive Mesh Refinement
AMT	Asynchronous-Many-Task
ANL	Argonne National Laboratory
APEX	Alliance for Application Performance at Extreme Scale
API	Application Programming Interface
ASC	Advanced Simulation and Computing (formerly ASCI)
ASCI	Accelerated Strategic Computing Initiative
ASCR	Office of Science's Advanced Scientific Computing Research
AT	Advanced Technology
ATCC	Advanced Technology Computing Campaign
ATDM	Advanced Technology Development and Mitigation
B&R	Budget and Reporting
CAD	Computer Aided Design
CBTF	Component-Based Tool Framework
CCC	Capability Computing Campaign
CCE	Common Computing Environment

CD	Critical Decision
CLAMR	Compute Language Adaptive Mesh Refinement
CoE	Center of Excellence
CORAL	Collaboration of Oak Ridge, Argonne, and Livermore
CPEP	Corporate Performance Evaluation Process
CPU	Central Processing Unit
CR	Continuing Resolution
CSSE	Computational Systems and Software Environment
CT	Commodity Technology
DOE	Department of Energy
DRAM	Dynamic Random Access Memory
DSL	Domain-Specific Language
DSW	Directed Stockpile Work
ECI	Exascale Computing Initiative
FOUS	Facility Operations and User Support
FPGA	Field Programmable Gate Array
FSEFI	Fine-Grained Soft Error Fault Injection Tool
GPFS	Global Parallel File System
GPGPU	General-Purpose Graphics Processing Units
HIO	High Performance Input/Output
HMC	Hybrid Memory Cube
HPC	High Performance Computing
HPSS	High-Performance Storage System
HQ	Headquarters
I/O	Input/Output
IB	Infiniband
IC	Integrated Codes
IDC	Integrated Design Codes
IDP	Identity Provider
IP	Implementation Plan
ITIL	Information Technology Infrastructure Library

LADAR	Laser Detection and Ranging
LAN	Large Area Network
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LDMS	Lightweight Distributed Metric Service
LEP	Life Extension Program
LFLR	Local Failure Local Recovery
LLNL	Lawrence Livermore National Laboratory
LLVM	Low-Level Virtual Machine
MOU	Memorandum of Understanding
MPI	Message Passing Interface
MST	Multisite Target
NAS	Network-Attached Storage
NFS	Network File System
NGBB	Next-Generation Backbone
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NRE	Non-Recurring Engineering
NSCC	National Security Computing Center
NVRAM	Non-Volatile Random Access Memory
nWBS	National Work Breakdown Structure
NWC	Nuclear Weapons Council
O|SS	Open|SpeedShop
OCF	Open Computing Facility
ORNL	Oak Ridge National Laboratory
OS	Operating System
PCF	Predictive Capability Framework
PDE	Partial Differential Equation
PECASE	Presidential Early Career Award for Scientists and Engineers
PEM	Physics and Engineering Models

PI	Principal Investigator
PIV	Personal Identity Verification
PKI	Public Key Infrastructure
PLFS	Parallel Log File System
PSAAP	Predictive Science Academic Alliance Program
PSP	Predictive Science Panel
QASPR	Qualification Alternatives to the Sandia Pulsed Reactor
QMU	Quantification of Margins and Uncertainties
R&D	Research and Development
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
RTS	Runtime System
SAN	Storage Area Network
SARAPE	Synchronized Account Request Automated Process
SC	Department of Energy's Office of Science
SCC	Strategic Computing Complex at Los Alamos
SCF	Secure Computing Facility
SCIDAC	Scientific Discovery through Advanced Computing
SCN	Sandia Classified Network
SFI	Significant Finding Investigation
SIO	Scalable Input/Output
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SRAM	Static Random Access Memory
SRN	Sandia Restricted Network
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
SUPER	SUstained Performance, Energy & Resilience Institutes
TLCC	Tri-Lab Linux Capacity Cluster
TOSS	Tripod Operating System Software
UGT	Underground Nuclear Test

UQ	Uncertainty Quantification
V&V	Verification and Validation
VM	Virtual Machine
VTK	Visualization Toolkit
WAN	Wide Area Network
WBS	Work Breakdown Structure
WCI	Weapons & Complex Integration (at LLNL)

Appendix B: Integrated Codes Subprogram (WBS 1.2.3.1)

Note: The content for the IC subprogram is available upon request from the ASC Program Office.

Appendix C: Physics and Engineering Models Subprogram (WBS 1.2.3.2)

Note: The content for the PEM subprogram is available upon request from the ASC Program Office.

Appendix D: Verification and Validation Subprogram (WBS 1.2.3.3)

Note: The content for the V&V subprogram is available upon request from the ASC Program Office.

Appendix E: Advanced Technology Development and Mitigation Subprogram (WBS 1.2.3.4)

Note: Additional content for the ATDM subprogram is available upon request from the ASC Program Office.

The ATDM subprogram includes laboratory code and computer engineering and science projects that pursue long-term simulation and computing goals relevant to the broad national security missions of the NNSA. It addresses the need to adapt current integrated design codes and build new codes that are attuned to emerging computing technologies. Performing this work within the scope of the DOE Exascale Computing Initiative (ECI) allows for broader engagement in co-design activities and provides a conduit to HPC vendors to enable next-generation, advanced computing technologies to be of service to the stockpile stewardship mission. Applications developers, along with computational and computer scientists are to build a computational infrastructure and develop a new generation of weapon design codes that will efficiently utilize the hardware capabilities anticipated in exascale-class systems.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2014, and through quarter 3, fiscal year 2015, are reflected below for the ATDM subprogram.

- Created three new teams: AAPS, production-quality tools development, and workflow; defined collaboration plan with IC, CSSE, and ASC codes that supports ATDM application efforts (LLNL)
- Conducted co-design activities and participated in the technical management of DesignForward projects (LLNL)
- Executed LLNL portion of the Level 2 milestone to demonstrate advances in proxy applications through programming abstractions or performance gains (LLNL)
- Conducted research and development (R&D) activities related to effectively utilizing next-generation architectures, including performance modeling, debugging at scale, code correctness on-node and at scale, power-aware HPC, support for fault tolerance in the context of current models, and novel scheduling concepts (elasticity and I/O-aware scheduling) (LLNL)
- Provided technical coordination and contractual management for FastForward contracts (LLNL)
- Completed Legion programming model investigation for FY14 Level 2 milestone (LANL)

- Executed and documented a requirements gathering exercise via a task force comprised of technical staff from across the NNSA weapons programs and ASC Program elements (LANL)
- Completed proxy application performance on emerging hardware FY14 Level 2 milestone (LANL)
- Defined initial design of execution, data, and state models for ATDM/next-generation code (LANL)
- Demonstrated initial software engineering environment for ATDM/ next-generation code (LANL)
- Executed ATDM Level 2 milestone "*Programming Models Analysis for Next-Generation Platforms*" by prototyping asynchronous task-based programming models on Mantevo mini-applications (SNL)
- Developed implementations of miniAero in Legion, Charm++, and Uintah that will be added to mantevo.org and used in the comparative study for SNL's FY15 Level 2 milestone (SNL)
- Delivered a five-year plan for deployment of the Kokkos performance portability library to ATDM codes in collaboration with the task parallel and data management activities (SNL)
- Developed initial software architecture for the data-warehouse service; captured requirements from meshing, particle, and multitask teams; results of these comparisons are driving the design and implementation of an updated prototype (SNL)
- Deployed ATDM collaboration environment including a SharePoint document repository, a Git software repository with individual workgroup access controls, a Trac issue tracking system, and a Jenkins continuous integration testing system (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Deliver Strategic Plan and Initial Scalability Assessment by Advanced Architecture and Portability Specialists Team		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: ATDM		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The newly formed advanced architecture and portability specialists (AAPS) team will develop a strategic plan to meet the goals of 1) sharing knowledge and experience with code teams to ensure that ASC codes run well on new architectures, and 2) supplying skilled computational scientists to put the strategy into practice. The plan will be delivered to ASC management in the first quarter. By the fourth quarter, the team will identify their first customers within PEM and IC, perform an initial assessment and scalability and performance bottleneck for next-generation architectures, and embed AAPS team members with customer code teams to assist with initial portability development within standalone kernels or proxy applications.		
Completion Criteria: When the strategic plan is signed off by ATDM management, the initial assessment is documented in an internal repository or report, and team members are embedded into at least two ASC code teams.		
Customer: Code teams with products that support next-generation code team efforts		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The “handoff” of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: IC, CSSE, Sequoia		

Milestone (ID#TBD): Next-Generation Code: Plan for Rapid Development		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: ATDM, IC, CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: The Next-Generation Code project has been conducting an exploratory and scoping phase (phase 1) involving evaluation of physics and computing methods and application (end-user) requirements. A toolkit of production-quality physics kernels, remap and link tools, and computer science methods, together with physics manuals and documentation, will be developed. A detailed plan will be developed for the rapid development of next-generation integrated codes in FY17–19 (phase 2), based on a down-selection or prioritization of the methods and application spaces evaluated in phase 1. The plan will include a functional design, definitions of metrics, and enumeration of risks to guide and evaluate progress. Integral problems will be identified for out-year milestones.		
Completion Criteria: Milestone completion will comprise three deliverables: 1) a report detailing phase 1 findings on physics, computer science, and application challenges and requirements; 2) a toolkit of production-quality physics kernels, remap and link tools, and computer science methods, together with physics manuals and documentation, available in the project software repository; and 3) a detailed plan for phase 2 rapid development of next-generation integrated codes.		
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: ATDM Architecture and Software Development, CSSE, and IC projects		

Milestone (ID#TBD): Demonstrate Performance Portability and Embedded Analysis Infrastructure		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: ATDM		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Given the importance of both performance portability abstractions (Kokkos) and the coupled multiphysics and template infrastructure for embedded analysis (Sacado), this Level 2 milestone is designed to ensure that the initial versions (v0) of both the thermal-mechanical safety application and the electromagnetic-plasma hostile environment ATDM application prototypes have demonstrated capabilities in both these dimensions by the end of FY16. Successful completion of this milestone requires a demonstration of performance portability on multiple hardware types utilizing the CSSE advanced architecture test beds. Completion also requires that the template automatic propagation of information be demonstrated for initial versions of both SNL ATDM demonstration applications.		
Completion Criteria: This milestone will be completed with 1) documented demonstration of performance portability on v0 of ATDM applications on ASC test beds, including Nvidia, Intel Phi, and traditional central processing unit (CPU) architectures, and 2) documented demonstration of templated automatic propagation information for coupled multiphysics applications; for example, a transient sensitivity analysis or a rudimentary UQ ensemble propagation would satisfy this requirement.		
Customer: ATDM and IC		
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 advanced architecture test beds		

Milestone (ID#TBD): Demonstrate Integration of Performance Improvements Identified via Proxy Applications into Application Codes		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: ATDM, IC, CSSE		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: This milestone is a tri-lab deliverable supporting the ongoing co-design efforts in the program as well as the new ATDM activities. In FY15, a milestone was completed evaluating proxy applications in multiple programming models. This milestone focuses on taking these improvements and extracting the knowledge gained and/or code revisions back into applications. This investigation will be done by evaluating the improvements in performance and portability that can be gleaned from previous work on proxy applications through prior Level 2 milestones and DesignForward/FastForward investigations. The identified improvements would then be applied to either IC applications or to new code in ATDM and be evaluated for inclusion. Successful and unsuccessful attempts will be reported as lessons learned. Suggested modifications to proxy applications will be reported to reflect what was learned during the integration process to produce more representative proxy applications.		
Completion Criteria: This milestone will be completed when 1) improvements from proxy applications from each lab have been identified and evaluated for applicability in IC or ATDM codes, and 2) a report has been completed by the tri-labs detailing lessons learned—both successes and failures—in regards to evaluation of performance improvements implemented into application codes, including how the proxy applications are representative of the application codes and where they could be improved.		
Customer: IC, ATDM		
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 advanced architecture test beds, IC, ATDM		

Projects for the Next-Generation Architecture and Software Development Product (WBS 1.2.3.4.2)

The Next-Generation Architecture and Software Development product is focused on long-term computing technology research to influence the shift in computing technology to extreme, heterogeneous architectures and to mitigate its impact and advance its capabilities for ASC simulation codes. This work is done in tight coordination with the other ATDM product, Next-Generation Code Development and Application. Projects perform computer science research on specific programming model technologies and computer architecture features anticipated in exascale computing. Targeted at exascale-class systems, the software will be evaluated on interim advanced architecture test beds and production AT systems.

Next-Generation Computing Enablement and Co-Design (LLNL)

The Next Generation Computing Enablement and Co-Design efforts are preparing ASC for the next generation of advanced computing technologies beyond the current AT system projects. This project includes coordinating next-generation activities with IC at LLNL, within the tri-lab, and externally. The efforts support IC ATDM next-generation application efforts to prepare for new platforms and to adapt other codes to the expected new architectures. The software efforts include system level software, development tools, data analysis tools, parallel I/O, and programming models. On the hardware side, these efforts include tracking and collaborating on technology innovations. This effort includes interactions with the DOE's ASCR, vendors, and academia, including planning and technical coordination for vendor contracts. Team members will carry out investigations and co-design activities using test beds and existing technology, making use of proxy applications.

Accomplishments in FY15:

- Created three new teams: AAPS, production-quality tools development, and workflow; defined collaboration plan with IC, CSSE, and ASC codes that supports ATDM application efforts
- Conducted co-design activities and participated in the technical management of DesignForward projects
- Executed LLNL portion of the Level 2 milestone to demonstrate advances in proxy applications through programming abstractions or performance gains
- Conducted R&D activities related to effectively utilizing next-generation architectures, including performance modeling, debugging at scale, and novel scheduling concepts (elasticity and I/O-aware scheduling)

Planned Activities in FY16:

- Conduct co-design activities with vendors and actively participate in Center of Excellence (CoE) efforts
- Participate in the technical management of DesignForward projects
- Conduct R&D activities in the tools area, including debugging at extreme scale
- Conduct R&D activities in the systems area, including next-generation resource management with novel scheduling algorithms and network congestion analysis
- Continue and expand investigation of new programming models and their runtime systems, especially task-based models, and collaborate closely with code team, next-generation efforts, and the AAPS team

Advanced Architectures and Portability Specialists Project (LLNL)

The LLNL ATDM strategy for supporting the development of next-generation simulation capabilities includes the formation of a team of specialists (AAPS) to facilitate the transfer and dissemination of hands-on advanced architecture expertise to code teams across ASC, rather than relying upon each code team to independently stay abreast of the latest developments in architecture, programming models and kernel optimization needed to make efficient use of new hardware. By staffing the team with experienced LLNL computational and computer scientists with a proven track record of success in scaling scientific applications on new, cutting-edge hardware, as well as hiring additional specialists in key areas such as general-purpose graphics processing unit (GPGPU) programming (CUDA/OpenACC or OpenMP 4), many-core programming, and parallel application development, LLNL will assemble a unique repository of knowledge and skills. This will both allow for the efficient dissemination of best practices as well as create an agile labor pool that can take a hands-on role in applying new approaches to existing codes.

Accomplishments in FY15:

- Created the team, hired the leader, moved next-generation developers onto the team and hired additional computational scientists to form the initial team
- Coordinated with LLNL IC code teams and CoE leaders to define mission and goals
- Developed plan for embedding team members with targeted code teams to gain in-depth experience and strengthen communication between the teams, as one part of the AAPS approach

Planned Activities in FY16:

- Stay abreast of new programming models developments taking place in the research community (for example, Predictive Science Academic Alliance Program (PSAAP II), ASCR X-stack program) and their usage in emerging next-generation proxy applications and production applications; perform independent evaluations on

promising technologies, and brief application teams on current status and LLNL assessment of readiness and applicability to LLNL codes

- Assess scalability and performance bottlenecks for next-generation architectures for multiple codes within IC and PEM, working closely with code teams
- Assist code teams with portability development within standalone kernels or proxy applications, including ALE3D and Cheetah
- Create repository of information about next-generation architecture knowledge, tools, and expertise, and make it accessible to the ASC community

Production-Quality Tools Development Project (LLNL)

Through a wide range of efforts as part of CSSE as well as leveraging ASCR-funded projects, LLNL has built up a strong portfolio in R&D of program development tools for performance analysis, debugging, correctness verification, power-aware computing, and resilience support. Further, several investments have been made in infrastructure projects in these areas. Enabling ASC code teams, especially in the context of ATDM efforts, to leverage and efficiently deploy software developed as part of these efforts is essential but also requires significant efforts to ensure the software is hardened to production quality, maintained on ASC systems, and documented for end users. The newly established production quality tools development team takes on the responsibilities to ensure these goals are met. It will work closely with the R&D efforts on one side and user support and AAPS teams on the other.

Accomplishments in FY15:

- Created the initial team, appointed a leader, and hired additional computer scientists
- Worked with ASC management to define mission and goals
- Developed plan for collaborating with AAPS team members, who will serve as expert tools users, beta testers, and collaborators with customers

Planned Activities in FY16:

- Develop strategic plan to meet the goal of producing production-quality software that is robust and easy to use, and delivering the tools to end users and providing support
- Develop and implement a testing and release strategy for program development tools, including guidelines on needed documentation, repository workflows, and long-term support and maintenance
- Identify commonalities in infrastructure needs and develop plan to integrate infrastructures towards a common software stack
- Identify initial tool set, based on user feedback (in particular, the AAPS team), and transition these tools to production

User Workflow and Modernization (LLNL)

In FY13 and FY14, LLNL undertook initial studies to identify how advances in big data could apply to the requirements of the weapons program, primarily in support of the designer community. Three areas in particular, and the specific needs called out by Weapons & Complex Integration (WCI) users, form the basis for the User Workflow and Modernization project:

- Problem Setup, including document gathering, utilization of geometric and design databases/sets, and meshing
- Simulation Management, tying together a multiplicity of codes, pre- and post-processing tools, and complex sets of input and output data
- Post-Processing, enabling ease of use for in-situ analysis and building provenance and data management into the simulated diagnostics methodologies

LLNL's current workflow environment is largely manually tracked and automated through scripting together a large and disparate set of tools and applications. High-quality science and engineering is produced with the workflows, so any solution must enhance rather than disrupt them. The goal is to optimize user productivity through the development of an advanced, integrated workflow environment at LLNL.

Accomplishments in FY15

- Assembled initial team, producing team goals and mission statements
- Organized live demonstration and initial interactions between WCI code teams and the Sandia Analysis Workbench project to determine whether its technologies could be useful for LLNL workflows
- Participated in user working group for workflow planning for next-generation tools
- Explored a Web application for Arbitrary Lagrangian-Eulerian (ALE) mesh management workflows combining job submission, visualization, and rollback

Planned Activities in FY16

- Articulate a long-term vision and get buy-in from stakeholders in WCI; the Applications, Simulations, Quality division; Livermore Computing (LC), and the Center for Applied Scientific Computing
- Identify and assess existing workflow tools and products (for example, geometry databases and representations, EOSView, and related tools)
- Build a team, including hiring at least one developer and bringing on existing staff to work on initial high-priority efforts
- Work with WCI users and developers to the identify highest priority areas for focus of new development

- Assess progress of existing programmatic and institutional projects in scientific workflows and determine where additional effort can be applied to bring new capabilities to users
- Explore the use of Web services as a possible methodology for tying together disparate workflow products
- Complete a prototype Web-based mesh management workflow application for KULL

FastForward—Industrial Partnerships for Extreme-Scale Technology Research and Development (LLNL)

The FastForward program is a jointly funded collaboration between DOE SC and NNSA to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme scale computing, on the path toward exascale computing. This program is administered by DOE and contracted through Lawrence Livermore National Security, LLC, as part of a seven-national laboratory consortium (Argonne (ANL), Lawrence Berkeley (LBNL), LLNL, LANL, Oak Ridge (ORNL), Pacific Northwest, and SNL).

The FastForward2 projects were awarded in late summer of 2014 and continue through calendar year 2016. The six projects address node-level architectural issues for exascale-class systems. Their specific focus areas are listed below:

- Advanced Micro Devices (AMD) Advanced Research has two projects: 1) study of multiple new memory architectures (including software support) and developing a processing-in-memory test bed and 2) integration and evaluation (including simulation) of node architectures considering processor designs, energy utilization, resilience, data movement, and programmability
- IBM Corporation has one project: flexible future memory interfaces and memory power efficiency
- Intel Federal LLC has one project: prototyping processor and node board designs within the context of the software stack and applications from the co-design centers
- NVIDIA Corporation has one project: node architecture investigations in energy, resilience, circuits and VLSI, network-in-chip, algorithms, and programmability
- Cray has one project: node architecture investigations in processors, memory, power, resiliency, compilers and runtimes

Accomplishments in FY15:

- Provided technical coordination and contractual management for FastForward contracts

Planned Activities in FY16:

- Provide technical coordination and contractual management for FastForward2

Cross-Cutting Extreme-Scale Research (LANL)

This is a cross-cutting applied research project that spans the program to address the challenges of new hardware and software at extreme scales. This includes research in advanced file systems, architectural analysis, workflows, and resilience and the systems interface for these. This work will be the majority of the interfaces from DesignForward and FastForward and other vendor interactions to the rest of ATDM and CSSE. The work will be done with the related CSSE projects to leverage and extend the expertise.

Accomplishments in FY15:

This is a new project beginning in FY16.

Planned Activities in FY16:

- Improve the data management, workflow, and analytics with novel investigations in data-management and workflow automation
- Define the long-term storage architecture in support of advanced applications on future platforms
- Begin an analytical analysis of applications and validate these applications on cutting-edge hardware and extreme-scale systems
- Expand the resilience analysis capability for advanced applications on future extreme-scale systems

Co-Design and Programming Model Research (LANL)

This project contains the forward-looking research for advance computing technologies at extreme-scale. Co-design research and programming model research are the base of these investigations in support of ASC code needs on future hardware.

The co-design component of the project leverages other activities at LANL to build a co-design process through the collaborative creation of patterns, strategies, and abstractions for the implementation and optimization of scientific applications and algorithms on emerging hardware architectures. One aspect of this process will be a suite of open-source proxy applications, derived from and feeding back into ASC IC teams. ASC code teams have informed and continue to provide the requirements for the study based upon application domains of interest to ASC. Specification documents and reference implementations produced will act as the basis for most of the work in this project, ensuring that it is targeted directly to ASC code needs.

The objective of this project is to track the development of next-generation hardware architectures and study both computational and data movement patterns represented by the chosen proxy applications developed in conjunction with ASC IC/ATDM code developers. A major goal of this effort is to inform application developers of methods and best practices that will be necessary for code development on AT system architectures such as Trinity. Work being done directly with IC/ATDM teams will support the evolution of current codes towards next-generation architectures by providing computer

science expertise on improved mesh data structures, new strategies for adaptive mesh refinement (AMR), performance improvement, data locality, and compressed data structures for materials. In addition, LANL will explore increased software abstraction, through emerging programming models and domain-specific languages (DSLs).

The programming models aspect of the project studies emerging hardware and software trends and their impact on programming abstractions/models. This includes the overall software development tool chain and run-time systems support for emerging programming models. LANL's goal is to develop a set of tools and technologies that will assist in the development of the next generation of application codes as well as extend the lifetime of current codes at extreme-scale and with vast parallelism.

In FY16, this project will be responsible for a tri-lab Level 2 milestone with IC/ATDM that will investigate the integration of knowledge gained from proxy applications back into production codes.

Accomplishments in FY15:

- Investigated alternate programming models with proxy applications, for a tri-lab Level 2 milestone
- Developed C++ (multi-core, GPU, Legion) framework for applicability to SNAP proxy application
- Developed a Kokkos array implementation of inner loop of SNAP
- Recorded success with the QUO thread arbitration library that manages multiphysics applications with message passing interface (MPI) and MPI+OpenMP, which showed runtime reductions of 10–20%
- Completed a noise study of the impact on tasking models versus MPI, MPI+OpenACC

Planned Activities in FY16:

- Further develop proxy application for ATS-*, DesignForward, FastForward interactions and identify improvements to IC codes as part of tri-lab Level 2 milestone
- Develop common physics abstractions for modern programming models
- Develop tools for extreme-scale systems, including the development of Byfl: a hardware-independent performance analysis and task-aware scalable debugger
- Investigate persistent object interface for task/data parallel programming models
- Develop low-level interfaces using low-level virtual machine (LLVM) for emerging programming models

Programming Models and Abstractions (SNL)

The next generation of computing platforms promises both new capabilities and increased capacity for meeting SNL's mission challenges. However, these platforms involve new computer architectures and using this new capability effectively requires an investment in new codes and algorithms that can effectively express and exploit multiple levels of parallelism. This project explores programing models and abstractions that will allow the new ATDM codes to effectively utilize next-generation hardware accounting for hybrid parallelism and the need for enhanced resilience. As part of the risk mitigation ATDM strategy, SNL will also invest and mature transitional models such as Kokkos that not only provide a means of implementing new codes, but can also be used to refactor low-level kernels in existing IDCs. To mitigate unknown characteristics of future computer architectures, SNL is developing asynchronous multitask programming models with an associated data-warehouse capability.

Accomplishments in FY15:

- Executed SNL's portion of the tri-lab Level 2 milestone *Demonstrate Advances in Proxy Applications Through Performance Gains and/or Performance Portable Abstractions*
- Executed ATDM Level 2 milestone *Programming Models Analysis for Next-Generation Platforms* by prototyping asynchronous task-based programming models on Mantevo mini-applications
- Delivered a five-year plan for deployment of Kokkos to ATDM codes in collaboration with the task parallel and data management activities
- Developed implementations of miniAero in Legion, Charm++, and Uintah (in collaboration with Next-Generation Code Development and Application Product) that will be added to mantevo.org and is being used in the comparative study for SNL's FY15 Level 2 milestone
- Began the design and documentation of an asynchronous-many-task (AMT) programming model and runtime systems (RTS) specification that outlines the requirements and application programming interfaces (APIs) for an AMT RTS based on SNL's ATDM workload
- Hosted three deep-dive boot camps on AMT RTS (Uintah, Legion, and Charm++), bringing ATDM applications and computer science infrastructure teams together with experts from each AMT RTS to develop implementations of miniAero for SNL's FY15 Level 2 milestone
- Developed initial software architecture for the data-warehouse service; captured requirements from meshing, particle, and multitask teams (results of these comparisons are driving the design and implementation of an updated prototype)
- Evaluated suitability and performance of alternate communication libraries that could be utilized as the backbone for the data warehouse (experiments revealed that the

Network Scalable Service InterfacE (Nessie) provides comparable performance and portability and offers better dynamic memory management opportunities)

- Hosted a Supercomputing 2014 Birds of a Feather on AMT Programming Models for next-generation HPC platforms to develop a conduit to the greater research community to benefit ATDM
- Collaborated with four undergraduate students from Harvey Mudd College to re-engineer selected computational kernels for thread-scalable parallelism using Kokkos
- Held a workshop on AMT RTS for next-generation HPC platforms on November 5–6, 2014, that helped develop a shared vision of the requirements and implementation plan for an AMT RTS for next-generation platforms
- Designed and implemented the DHARMA transport layer, a lightweight communication transport layer that provides fault-tolerant collectives for global agreement, flexible active messages via C++ object migration, well-defined reliability semantics for point-to-point communication, and asynchronous progress threads

Planned Activities in FY16:

- Continue enhanced productization of Kokkos to enable accelerated maturation, deployment, and support for Kokkos as a production capability to both ATDM and current production codes
- Deliver a Kokkos user programming guide and hold in-depth tutorial(s) for ATDM code developers
- Deliver Kokkos capability in support of the ATDM Level 2 milestone, including evaluation of performance of embedded analysis components (Sacado) and performance gaps
- Develop Kokkos back-end for the Trinity architecture using the Knights Landing test bed and evaluate performance
- Initiate development of a *Kokkos Kernels* library of cross-cutting (domain non-specific) algorithms and data structures built on Kokkos that can be used as initial templates and tutorials for algorithms and applications developers
- Improve performance of ASC application codes using Kokkos in support of the tri-lab co-design Level 2 milestone; document lessons learned
- Provide an initial AMT API to the ATDM applications, solvers, and analysis teams
- Provide an initial implementation of an AMT runtime capable of supporting code development; deliver a report summarizing lessons learned, best practices, capability gaps, and continued research directions based on feedback from application, solver, and analysis teams
- Provide support for AMT interoperability with Kokkos and explore the effects of task-granularity and over-decomposition (joint with Kokkos and ATDM application teams)

- Implement and demonstrate a key-value storage service that executes a separate application to allow sharing between coupled applications and application workflow components
- Implement a prototype data management abstraction layer that supports meshes, fields, and particle data as well as a key-value system that supports non-volatile storage (for example, burst buffer)
- Mature and harden on-node runtime research (for example, Qthreads) to enhance vertical integration across ATDM applications, Kokkos, and the lower-level exascale runtime (in coordination with CSSE)
- Integrate on-node RTS (Qthreads) with ATDM Kokkos, AMT, and data warehouse components (in collaboration with CSSE)

AgileComponents (SNL)

SNL's ATDM Strategy strongly leverages the AgileComponents software infrastructure for building next-generation engineering software that emphasizes: agile and scalable development; modularity for sustainment and to reduce the cost to transform to production; agility with-respect-to changing architectures and use-cases; and extensibility to meet new requirements.

Sandia has developed the AgileComponent approach software deployed in the open-source libraries Trilinos, Dakota, and demonstrated across a broad range of application demonstrations that point the way to flexible, interoperable yet independent libraries with clearly defined and extensible APIs. Within ATDM, we will leverage and extend this approach to develop the next generation of AgileComponents to meet the needs of future ASC production applications while also adding support for advanced ATDM capabilities including Kokkos performance portability, asynchronous multi-task, the data warehouse, and embedded workflows in close collaboration with algorithms and applications developers. This project focuses on the lower level components and their APIs that are written at a degree of abstraction that can be used across a wide range of ATDM applications. These components often interface more directly with Kokkos, AMT and data warehouse, providing a means of insulation to more application specific component layers.

Accomplishments in FY15:

- Began initial design API designs for time-integration, physics coupling, and parallel-mesh database
- Completed the merge of Kokkos into the core AgileComponents (Phalanx, Intrepid, Panzer), laying the foundation for demonstration of performance portability in an FY16 Level 2 Milestone

Planned Activities in FY16:

- Deliver initial mesh database API; implement initial parallel unstructured mesh database component that provides basic topology information built on Kokkos parallel structures; demonstrate basic embedded meshing capability on a set of test problems designed to illustrate the workflow; demonstrate portability by documenting successful test problem results on ASC test beds, including Nvidia, Intel Phi, and traditional CPU architectures
- Deliver an initial time integration API that includes support for IMEX and adjoint sensitivity analysis; implement basic time integration methods needed to support ATDM applications for the Level 2 FY16 ATDM milestone; demonstrate temporal order of accuracy on basic physics test problems for the milestone on ASC test beds
- Demonstrate initial contact global and local search and imprinting component capabilities with a set of tests that show capabilities for both implicit and explicit time integration strategies (capability will focus on MPI+Kokkos parallelism)
- Support the Level 2 FY16 ATDM milestone by integrating Kokkos-based assembly capabilities (Sacado, Intrepid, Phalanx and Panzer) into ATDM applications; demonstrate assembly tool portability across ASC test beds
- Produce a document that outlines the overall software design and component interactions for the next-generation AgileComponents
- Produce a plan for the generalization of the AgileComponent solver and analysis technology that extends abstraction layers to solver technology interfaces; redesign the underlying software interface for Kokkos support
- Initiate R&D on advanced workflow models that leverage next-generation capabilities such as embedded meshing, geometry, optimization, and UQ within the Next-Generation Code Development product; and document approaches to machine learning and other data analytics technologies to make user-adaptive workflows

Software Development Infrastructure and Test Beds (SNL)

The SNL ATDM approach heavily leverages a component-based software design that presents several challenges to the software development team, including heavy templating, complex interactions, and a large number of software combinations that require testing. This project focuses on procurement, maintenance, and support of a dedicated ATDM development system that will be specifically provisioned to enable:

- Rapid builds of complex, template software with large shared memory nodes and very fast local I/O to stream through large amounts of header files
- A moderate number of interconnected compute nodes that will act as a test farm, enabling continuous integration testing that will allow SNL to test the ATDM software components and applications.

These compute nodes will be provisioned with both NVIDIA GPUs and Intel Phi accelerators, and future hardware as it becomes available, so that SNL can explicitly test performance portability and AMT implementations in heterogeneous environments. Note that this is complementary and very distinct from the current CSSE test beds, which are advanced, first-of-a-kind CS research platforms. Instead, this ATDM test bed is a stable hardware and software environment managed specifically to help expedite the development, testing, and productization of ATDM software. This ATDM development hardware and software infrastructure will be refreshed (considered on an annual cycle) but using mature versions of hardware and software, whereas CSSE test beds are on the cutting edge.

Accomplishments in FY15:

- Developed initial requirements and design for the ATDM development test bed and prepared for procurement in fourth quarter
- Deployed initial ATDM collaboration environment (in collaboration with the Next-Generation Code Development Product) including a SharePoint document repository, a Git software repository with individual workgroup access controls, a Trac issue tracking system, and a Jenkins continuous integration testing system

Planned Activities in FY16:

- Procure updates to the original ATDM development test bed deployed in FY15 to include advanced hardware such as updates to Nvidia and Intel Phi accelerators; provide support and maintenance for the ATDM development test bed
- Deploy software tools and capabilities such as integrated development environments, continuous integration testing, and debugging in support of ATDM software development; provide support to the ATDM software teams
- Provide an automated build system for third party libraries and version-of-the-day application components for development and test bed environments; deploy to ATDM software developers
- Develop a comprehensive plan for software tool support for ATDM applications; develop checklist(s) for adding new people to the ATDM project; document current practices used by the applications teams, including build tools, development workflow, testing frequency, and reporting; implement nightly and continuous integration testing for application teams

DesignForward—Industrial Partnerships for Extreme-Scale Technology Research and Development (LBNL)

The DesignForward program is a jointly funded collaboration between DOE SC and NNSA to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme scale computing, on the path toward exascale computing. This program is administered by DOE and contracted through LBNL, as part of a seven-lab consortium (ANL, LBNL, LLNL, LANL, ORNL, Pacific Northwest

National Laboratory, and SNL). DesignForward seeks to fund innovative new and/or accelerated R&D of technologies targeted for productization in the 5–10 year timeframe.

The first set of DesignForward contracts focused on network interconnect projects and was completed in FY15. The second set of DesignForward contracts focus on system integration and were awarded that same year and will complete in 2017.

Accomplishments in FY15:

- Successfully completed the AMD, Cray, IBM, Intel, and NVIDIA DesignForward Interconnect projects
- Awarded DesignForward (DF2) System Integration contracts:
    - AMD Advanced Research: R&D in a conceptual system design, execution models, metrics and their evaluation, and participation in co-design activities
    - Cray: studies in system components for an exascale system, an execution model runtime specification for two disparate programming models, and an advantage/disadvantage analysis of the studied system components
    - IBM: definition of an execution model, exploration of memory interfaces and architectures, performance estimation in different hardware configurations, definition of metrics for system evaluation, and consideration of design options based on workflows and proxy applications
    - Intel: creation of a multilayer refinement and abstraction workload tool for simulation of components, system performance, programming models, and system architecture tradeoffs
- Provided technical coordination and contractual management for both the Interconnect and System Integration DesignForward contracts

Planned Activities in FY16:

- Provide technical coordination and contractual management for the DesignForward contracts in system integration.

Appendix F: Computational Systems and Software Environment Subprogram (1.2.3.5)

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of the NNSA. This sub-program strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms. 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. Achieving this deployment objective requires sustained investment in applied R&D activities to create technologies that address ASC's unique mission-driven needs for scalability, parallelism, performance, and reliability.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2014, and through quarter 3, fiscal year 2015, are reflected below for the CSSE subprogram.

- Awarded the Sierra contract and began application preparations through the CoE (LLNL)
- Led the CTS-1 process, including releasing CTS-1 request for proposal (RFP), reviewing proposals, selecting vendor, and awarding the contract (LLNL)
- Released Tri-Lab Operating System Software (TOSS) 2.3 (LLNL)
- Ran two Capability Computing Campaign (CCC) processes on Sequoia (LLNL)
- Conducted software plan surveys for Sequoia and Sierra on 1) programming models, 2) code correctness, debugging, and performance analysis tools, and 3) power-limited computing and resilience issues; and documented findings (LLNL)
- Released an open source version of PINION, a portable, data-parallel software framework for physics simulations (LANL)
- Investigated fault injection/recovery with CLAMR hydro application (LANL)
- Enabled production runs of the Eulerian Applications Project codes on Sequoia (LANL)
- Demonstrated HIO, the burst buffer abstraction library, on test beds (LANL)

- Showed initial full-scale machine simulation models of key ASC computing platforms using SNL's Structural Simulation Toolkit (SST), pointing the way to accurate, scale-free performance predictions for important parallel workload algorithms (SNL)
- Released Version 1.0 of the PowerAPI prototype implementation, which is now available on http://powerapi.sandia.gov (SNL)
- Completed a study that characterizes the power sensitivity of mini-applications on multiple test bed architectures, analyzing both power states (fixed performance levels, varying power) and power caps (fixed power cap, variable performance) (SNL)
- Developed a node-aware data redundancy scheme for Local Failure Local Recovery (LFLR), a scalable resilience programing model (SNL)
- Developed ensemble analysis tools to support model tuning through comparison of simulation results with experimental data (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Modernization and Expansion of LLNL Archive Disk Cache		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 3/31/16		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Configuration of archival disk cache systems will be modernized to reduce fragmentation, and new, higher capacity disk subsystems will be deployed. This will enhance archival disk cache capability for ASC archive users, enabling files written to the archives to remain resident on disk for many (6–12) months, regardless of file size.		
Completion Criteria: Archival disk cache systems are modernized, and new, higher capacity disk subsystems are deployed.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: LLNL FOUS staff		

Milestone (ID#TBD): Deploy Commodity Technology System-1 Tripod Operating System Software Test Bed Scalable Unit (Hype)		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 6/30/16		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Installation and configuration of a single scalable unit of CTS-1 will be installed in the LLNL test bed. The system will be used to debug TOSS 3.0 on CT system hardware.		
Completion Criteria: Perform hardware and software integration and testing in the LLNL test bed of one scalable unit of the next generation of the CT systems.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The “handoff” of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: LLNL FOUS staff		

Milestone (ID#TBD): Trinity-Haswell High Performance Computing System Delivery		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 12/31/15		
ASC nWBS Subprogram: CSSE, FOUS		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: Prepare Trinity-Haswell subsystem for integration into the LANL Computing Center. The Intel Haswell processors are the currently available commodity chips deployed for HPC. The second half of the Trinity system will be delivered with the Intel Knights Landing processors in Q3FY16. Deliver and install system hardware. Deliver, test, and demonstrate system software. Complete onsite capability scaling testing. Prepare the system for onsite 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 onsite to the extent that is contractually required is substantially complete; 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 onsite capability scaling testing and demonstration; and system is ready to begin onsite 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, New Mexico Alliance for Computing at Extreme Scale (ACES), LANL facilities		

Milestone (ID#TBD): Demonstrate and Evaluate Advanced Analysis, Visualization, and Input/Output Capabilities for the SIERRA Toolkit		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 6/30/16		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: The combination of increased complexity of ASC simulations and projected limitations in I/O system performance for the next generation of extreme-scale systems is causing a significant change in the usage model for HPC systems. The standard approach of executing each step in the application workflow in distinct steps, storing intermediate results to a shared parallel file system, will not be practical. To address this challenge, CSSE has been developing technologies to enable a tighter coupling between different elements (for example, simulation and analysis) in the application workflow and it has been adapting analysis and visualization codes to execute on HPC systems alongside the large-scale simulations. This milestone will provide a demonstration and evaluation of the use of these capabilities for important applications using the SIERRA toolkit.		
Completion Criteria: Demonstrate in situ, in transit, and ensemble analysis on a weapons-relevant problem. Dakota will be used to manage the ensemble runs, Catalyst for in situ analysis and data formatting, and Slycat for the ensemble analysis.		
Customer: Nuclear weapons engineering analysts, SIERRA users		
Milestone Certification Method: A program review is conducted and its results are documented. The “handoff” of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: Support from SIERRA users to help demonstrate and evaluate this capability		

Milestone (ID#TBD): Evaluate Impact of Advanced Memory Architectures on ASC Codes		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Analyze the feasibility of next-generation memory system architectures to increase performance of ASC applications in partnership with industry and academia. SNL is performing architectural analysis and supplying application and system software expertise. Analysis of this system will be focused on quantifying improvements in memory system performance or power compared to conventional memory systems. Architectural simulation experiments will also be used to explore the memory system design space (including bandwidth, capacity, and topology); determine the overall impact on the applications, system software, and system balance; and determine the impact of performing some computation, synchronization, or data movement operations in the memory system. This milestone will focus on the performance of key ASC application kernels and algorithms and will use a variety of tools available at the time, which may include the SST simulator, hardware emulation prototypes such as field programmable gate arrays (FPGAs), or hardware evaluation test beds. Next-generation memory architectures include Micron's Hybrid Memory Cube (HMC) and possibly other memory architectures such as multilevel memory and other conceptual designs from the DOE FastForward R&D projects.		
Completion Criteria: A study report on the impact of different uses of a relevant next-generation memory system for a set of representative mini-applications or proxy applications representing ASC code kernels and algorithms.		
Customer: Nuclear weapons engineering code teams, to guide implementation of algorithms for future ASC platforms		
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: ASC advanced architecture test beds		

Projects for the Commodity Technology Systems Product (WBS 1.2.3.5.1)

The CT Systems product provides 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, procurement and integration coordination, and installation. This product also provides market research for future CT systems.

Production Planning and Integration (LLNL)

The LLNL ASC strategy for CT systems 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 commodity 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.

Accomplishments in FY15:

- Met with many vendors to review and influence product roadmaps for CTS-1 timeframe deliverables
- Released CTS-1 RFP, reviewed proposals, selected CTS-1 vendor, and awarded CTS-1 contract

Planned Activities in FY16:

- Extensively test and resolve issues on pre-production CTS-1 systems
- Deploy FY16 CTS-1 systems into LLNL simulation environments

Commodity Systems Planning and Deployment (LANL)

The scope of the Commodity Systems Planning and Deployment project is to support the design, acquisition, delivery, and deployment of CT production systems. Primary capabilities include the planning and coordination necessary to integrate, accept, and transition CT systems into the HPC production environment at LANL. Efforts include the development of design criteria based on LANL's ASC simulation workload and facility capability—as part of a tri-lab requirements planning team, support for the ASC CT

system acquisition strategy, and execution of the integration and stabilization activities of the CT systems.

Accomplishments in FY15:

- Continued to operate Luna and the other capacity systems in both the classified and unclassified computing environments
- Participated in the tri-lab selection for the NNSA ASC CTS-1 procurement

Planned Activities in FY16:

- Continue to provide production support for CT systems
- Deploy and integrate the initial CTS-1 systems into the classified and unclassified computing environments

ASC Commodity Systems (SNL)

The purpose of the ASC Commodity Systems project is to support the acquisition, delivery, and installation of new ASC CT systems. The project is supported by analysis of SNL's portfolio of application needs for capacity workload systems within the context of the higher integrated ASC platform strategy of commodity and AT systems. Efforts include definition of requirements for CT systems and collaboration with the CCE product, with respect to a common software stack for new and existing capacity systems.

Accomplishments in FY15:

- Supported tri-lab survey, RFP development, proposal technical and business review to complete vendor selection and initiation of procurement contract for CTS-1

Planned Activities in FY16:

- Prepare facility space and infrastructure and support installation, testing, facility and network integration, security, and production operations for CTS-1 system(s)

Projects for the Advanced Technology Systems Product (WBS 1.2.3.5.2)

The AT systems 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.

Sequoia Tri-Lab Advanced Technology Platform (LLNL)

Sequoia is a 20-petaFLOP/s IBM BlueGene/Q system platform that was sited at LLNL in FY12 with final acceptance in early FY13. BlueGene/Q brings many innovations over the previous BlueGene generations, including 16 cores per node, multithreaded cores, a five-dimensional torus interconnect, water cooling, and optical fiber links. The 20-petaFLOP/s system has a staggering 1.6 million processor cores with a total possible 102 million hardware threads all operating simultaneously. This type of parallelism dictates new directions in supercomputing and enters a new regime of the possible physical systems that can be simulated numerically. Codes that are optimized for multi-core and multi-threading will run best on this machine. This platform will be used as a CCC machine for tri-lab stockpile stewardship milestones. Every six months a new CCC process will be run and the next suite of codes will be ushered onto the machine.

Accomplishments in FY15:

- Ran two CCC processes
- Investigated optimal performance tuning

Planned Activities in FY16:

- Run two CCC processes
- Continue to investigate optimal performance tuning for specific codes

Sierra Tri-Lab Advanced Technology System (LLNL)

In November 2014, LLNL signed a contract with IBM to deliver a next-generation supercomputer in 2017. The system, to be called Sierra, will serve the NNSA ASC

Program. Procurement of Sierra is part of the Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) laboratories to accelerate the development of HPC. Under the contracts, LANL and ORNL will work with IBM, NVIDIA, and Mellanox to deploy systems of over 100 petaFLOPS/s to advance science and ensure national security. Sierra will be a key tool for the three NNSA laboratories in pursuing predictive applications necessary to sustain the nation's nuclear deterrent and dedicated to high-resolution weapons science and UQ for weapons assessment.

Accomplishments in FY15:

- Awarded the Sierra contract on November 4, 2014
- Provided technical coordination and contractual management for CORAL non-recurring engineering (NRE) and Sierra contracts
- Awarded the ANL CORAL NRE contract to Intel on March 9, 2015
- Began application preparations for Sierra system through the CoE

Planned Activities in FY16:

- Provide technical coordination and contractual management for CORAL NRE and Sierra contracts
- Participate in NRE and build subcontract activities for ANL's CORAL system, Aurora, as part of the broader CORAL collaboration activities
- Deploy the Sierra early access system to help applications prepare for Sierra

Alliance for Computing at Extreme Scale Trinity Advanced Technology System (LANL, SNL)

The objective of this project is to define requirements and potential system architectures for platforms that meet future ASC programmatic requirements and drivers. The primary activity is to lead the design, acquisition, and plan for deployment of the Trinity AT system. The project will take into consideration mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design and operation process.

The Trinity platform is the first AT system for the ASC Program per the *2013 Computing Strategy*[4]. The project is a joint collaboration of ACES, a partnership between Los Alamos National Laboratory and Sandia National Laboratories.

The architecture and design of Trinity is to provide performance for large-scale applications in support of the NNSA program's most challenging problems. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

[4] *ASC Computing Strategy*, 2013, issued by the Office of Advanced Simulation & Computing, NA-114, SAND 2013-3951P.

Trinity will replace the Cielo platform sited at LANL but will be used by the applications users at each of the NNSA labs.

Accomplishments in FY15:

- Installed and accepted the Trinity test bed systems
- Instantiated the CoE for Trinity
- Released Version 1.0 of the PowerAPI prototype implementation, which is available on powerapi.sandia.gov
- Held Trinity quarterly reviews with Cray in December and April
- Received delivery of initial Haswell-based Trinity test beds at LANL and SNL
- Installed the Trinity-Haswell system

Planned Activities in FY16:

- Complete the Level 2 milestone *Trinity-Haswell High Performance Computing System Delivery*
- Continue to provide technical coordination and management of the Trinity contract
- Complete acceptance of the Trinity-Haswell system
- Conduct Open Science runs on the Trinity-Haswell system
- Transition the Trinity-Haswell system into the classified environment after the open science runs are complete
- Run the initial Advanced Technology Computing Campaign (ATCC-1)
- Advance Trinity project, including initial capability for advanced power management and testing of burst buffer capabilities
- Work closely with Cray to ensure final delivery of Trinity-KNL part of the system (Xeon Phi processors) per planned schedule

Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (LANL, SNL)

The Cielo capability computing platform is a project under the ACES. ACES is a joint collaboration between LANL and SNL defined under an MOU 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.

The architecture and design of Cielo is optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs. This project

covers all aspects of the technical, programmatic, and procurement planning for the platform.

Cielo is the ACES platform sited at LANL that supports the ASC CCCs. Cielo provides 1.37 peak petaFLOP/s with over 140,000 compute cores, and 10 petabytes of storage. Over 6,000 of the cores are dedicated to visualization and data services activities with connections to the SNL and LLNL sites.

Accomplishments in FY15:

- Supported ongoing Cielo production work in support of CCCs
- Facilitated dedicated time within Cielo CCC-7 to support critical B61-12 LEP program deliverables for engineering design and assessment
- Completed Cielo CCC-7; completed CCC-8 proposal review and initiated campaign

Planned Activities in FY16

- Continue to run Cielo in production capability mode
- Support improved code performance, analyst workflows, and proposal process and workload prioritization for simulation on future CCCs
- Complete Cielo CCC-8; begin and complete CCC-9
- Decommission Cielo

Alliance for Computing at Extreme Scale Crossroads Advanced Technology System (LANL, SNL)

The objective of this project is to define requirements and potential system architectures for platforms that meet future ASC programmatic requirements and drivers in the 2021–2025 timeframe. The primary activity is to lead the design, acquisition, and deployment of the Crossroads system. The project will take into consideration mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design and operation process.

In keeping with the mission requirement to field an AT system, some portion of the Crossroads procurement budget will be devoted to NRE work in partnership with the selected vendor.

ACES is continuing the partnership with DOE/SC LBNL to acquire two AT systems in the 2020 timeframe (one to be sited at LBNL and one at LANL). This collaboration is called APEX (Alliance for Application Performance at Extreme Scale).

The architecture and design of Crossroads are to provide performance for large-scale applications in support of the NNSA program's most challenging problems. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

Crossroads will replace the Trinity system sited at LANL but will be used by the applications users at each of the NNSA labs.

Accomplishments in FY15:

- Completed the Critical Decision 0 (CD-0) Mission Needs Package process to initiate the project in FY16

Planned Activities in FY16

- Develop the acquisition strategy and RFP
- Complete the CD-1/3a Conceptual Baseline Package and obtain concurrence

Projects for the System Software and Tools Product (WBS 1.2.3.5.3)

This level 4 product provides the system software infrastructure, including the supporting operating system (OS) 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 Environment for Scalable Systems (LLNL)

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 OS 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 OS and interconnect (for example, TOSS and Infiniband (IB) monitoring tools), as well as the resource management environment (Moab and Simple Linux Utility for Resource Management (SLURM)) to queue and schedule code runs across LLNL systems. LLNL uses TOSS on all of its Linux clusters. 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 FOUS' TOSS funding accounts for 40 percent of the effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in section 1.5.5.6 of this document.

Accomplishments in FY15:

- Released updates to TOSS (version TOSS 2.2.1–1, 2.2.1–3, and 2.3–3) that included Lustre version 2.5.3, security updates, and bug fixes
- Released TOSS 2.3 (based on Red Hat Enterprise Linux (RHEL) 6.6, the latest release from Red Hat)
- Developed TOSS 3 (based on RHEL 7)

- Developed enhancements to the IB OpenSM Monitoring Service, which provides an application interface to IB diagnostics, monitoring, management, and control functions
- Investigated alternative architectures for commodity Linux clusters (specifically, ARM)

Planned Activities in FY16:

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 2.X for legacy systems (based on RHEL 6.X), with X defined at actual release time
- Develop/deploy TOSS 3.X for CTS-1 systems (based on RHEL 7.X), with X defined at actual release time
- Continue to develop enhancements to the IB OpenSM Monitoring Service, which provides an application interface to IB diagnostics, monitoring, management, and control functions
- Develop a prototype software architecture that can continuously collect and analyze system data produced by the Lightweight Distributed Metric Service (LDMS) along with performance measurements from WCI proxy applications
- Investigate alternative architectures for commodity Linux clusters (for example, ARM)

Applications Development Environment and Performance Team (LLNL)

The Applications Development Environment and Performance Team (ADEPT) project provides the code development environment for all major LLNL platforms, supports user productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The elements of the development environment covered by this project include, but are not limited to, compilers, debuggers, power and resilience, performance and memory tools, interfaces to the parallel environment, and associated run time library work.

Accomplishments in FY15:

- Completed Level 2 milestone report on post-petascale development environment, and participated in scoping Sierra development environment requirements
- Advanced the state of the Sequoia development environment; supported tri-lab code teams on performance and porting needs for the CCC
- Developed tool prototypes for correctness and verification efforts, both for shared memory and message passing environments

- Started prototyping a standardized debug interface for OpenMP programs
- Developed tools characterizing communication performance, both on-node and cross-node

Planned Activities in FY16:

- Conduct co-design activities with vendors (for example, complete prototype of the OpenMP debug interface and work with RogueWave team to integrate it into Totalview, as part of Sierra NRE)
- Continue support of CCC efforts, and when possible, enhance the Sequoia development environment to support emerging programming requirements for applications codes (examples are certain MPI3 and OpenMP capabilities being released and incorporated into production codes in preparation for future systems)
- Expand efforts on correctness and verification tools; transition existing prototypes to a state in which they can easily be used by code teams
- Conduct initial testing and provide long-term support of new CTS-1 environment
- Develop new performance visualization techniques, both building on top of the existing communication characterization tools and using new data sources, including power and resilience data; focus on novel correlation techniques to detect root causes of performance problems

High Performance Computing Systems Research (LANL)

HPC systems research is a broad-reaching project focusing on near- to long-term research of all the components needed to support a rich environment for very large-scale applications. Systems research bridges the gap between hardware and programming model, and requires tight collaboration in supporting the development of programming models, tools, visualization/analytics, and system software aspect of I/O.

The project includes investigations on resilient system services, soft-error resilience, system support for data-intensive computing, power, and interconnect topology modeling/evaluation.

Resilient system services focus on developing a vehicle to investigate resilient, dynamic, distributed, scalable services for large-scale systems and providing an interface to programming models so that ASC applications can access these features on current and future hardware. Current activities include investigation of distributed systems software for job launch and monitoring. An imminent challenge for extreme-scale HPC systems is the issue of power limits and rapidly varying demands on the grid. Techniques for power-capping HPC systems will be investigated.

Investigations of soft-error resilience will continue work on an accelerated testing environment for soft error profiling using a virtual machine (VM) approach to inject faults while actively running real ASC codes. Hardware reliability in HPC systems remains a challenge to characterize. Statistical studies of reliability data from a variety of

production systems will be extended, and models of dynamic random access memory (DRAM) reliability will be developed. Reliability of non-volatile storage will be studied in support of future hierarchical storage systems (for example, burst buffer architectures).

Accomplishments in FY15:

- Demonstrated fault injection and recovery with the Compute Language Adaptive Mesh Refinement (CLAMR) hydro application
- Developed Hop, a distributed key value store with per key consistency targeted for systems software and services

Planned Activities in FY16:

- Analyze and extend models of DOE production system statistics related to reliability, including system characterization, code characterization, and static random access memory (SRAM) aging
- Develop further Fine-Grained Soft Error Fault Injection Tool (FSEFI), fault injection, and fault analysis for DOE workloads
- Extend investigation of distributed key-value stores and I/O middleware
- Model interconnect performance and mixed scheduling of MPI and fine grained-ensemble workloads with data dependencies and power capping

Advanced System Test Beds (LANL)

The Advanced System Test Beds project provides program management for CSSE and provides test bed hardware and software for research investigations in support of the IC/ATDM/CSSE missions. It fills the gaps of advanced architecture hardware and provides local access to advanced hardware.

Accomplishments in FY15:

- Completed a technology refresh for Darwin, the advanced architecture test bed
- Completed an initial investigation of Micron HMC “pico card”

Planned Activities in FY16:

- Conduct performance investigations of new processors, including ARM, Power, and AMD, including their corresponding deep/complex memory hierarchies
- Continue minor rolling technology refresh for advanced architecture test bed

System Software Stack Advancement (SNL)

The System Software Stack Advancement project supports system software R&D to address scalability and efficiency of future computational systems in multiple dimensions. An important aspect is providing lightweight services and functionality that

does not compromise scalability and therefore performance. The focus will be on three critical areas for HPC systems, which will enhance efficiency, performance and scalability of applications on future HPC systems:

- Power has become a first-order design constraint for future supercomputers. SNL will expand upon work in data collection and tuning techniques that provided new insight into understanding power requirements and affecting power use of ASC applications.
- Much of the complexity of managing future heterogeneous compute and memory resources will be the responsibility of the RTS. SNL will continue to explore the relationship between the RTS, the OS, and the interconnect to provide the necessary policies and mechanisms for ensuring scalability and performance while insulating the complexities of the resources from applications.
- Previous work with virtualization has shown promise in the area of HPC. Virtualization will be leveraged to provide insights into application runtime characteristics and where optimization efforts would be best targeted.

As a long-term goal, SNL plans to integrate these targeted efforts with previous successes in lightweight OS (Kitten), lightweight RTS (Qthreads), and high performance network stack (Portals communication protocol) development with a production HPC computing stack. While this is a significant development effort, the goals are to provide the program the following:

- Risk mitigation against vendor provided software failure
- Setting the standard for vendor-delivered scalable software stacks
- More complete understanding of the responsibilities of system software relative to applications

This effort will necessarily be accomplished in conjunction with the acquisition of future AT systems.

Accomplishments in FY15:

- Completed a study that characterized the power sensitivity of mini-applications on multiple test bed architectures, analyzing both power states (fixed performance levels, varying power) and power caps (fixed power cap and variable performance)
- Completed a prototype of a subset of the PowerAPI specification using PowerInsight v2 as the measurement and control mechanism

Planned Activities in FY16:

- Evaluate the advanced power management capabilities of the Trinity platform for ASC workloads; work with Cray to enhance and extend where necessary
- Deploy PowerAPI prototype implementation on test bed platforms at SNL
- Deploy PowerAPI on CT systems via integration with TOSS

High Performance Computing Hardware Architecture Simulation (SNL)

The SST is a suite of tools enabling multi-scale computer architecture simulation to meet the needs of HPC software/hardware co-design. 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, permitting the consideration of a broad space of potential architectural and application/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.

Accomplishments in FY15:

- Developed capability to analyze Sandia HPC systems with coarse-grained application models running on high-fidelity interconnection network models
- Showed initial full-scale machine simulation models of key ASC computing platforms using SNL's SST, building on two years of vendor engagements through the DOE's FastForward and DesignForward programs, pointing the way to accurate, scale-free performance predictions for important parallel workload algorithms
- Hosted SST tutorial session at the International Symposium on Computer Architecture 2015, which was presented on June 14, 2015, following the Gem5 workshop on June 13, 2015 (these efforts are intended to support the growth of SST users and developers)

Planned Activities in FY16:

- Integrate SST-Micro and SST-Macro with a common repository and build system
- Assess impact of advanced memory architectures on ASC codes
- Model and analyze Trinity Dragonfly interconnection network performance to determine application improvements in communication and to optimize performance on Trinity

Interprocess Communication System Software Stack (SNL)

The Interprocess Communication System Software Stack project will develop capabilities to enable performance and scalability of ASC applications on current and future high-performance interconnection networks on extreme-scale platforms. This project will concentrate on characterizing application requirements with respect to functionality and

performance for intra-application data movement as well as application network transfers to external I/O services. It will also provide a low-level network programming interface appropriate for current-generation network hardware as well as more advanced next-generation hardware with more sophisticated network interface capabilities and functionality. As applications explore alternative programming models beyond the current distributed memory MPI model, the low-level network programming interface must evolve to include the ability to provide very lightweight one-sided data transfer operations, while continuing to enable efficient two-sided message-based transfers. It is likely that this project will expand to include an analysis of network topologies, network interface hardware design and evaluation, optimized network transfer protocols, and system software support for advanced network interface operations.

This project will build on existing efforts surrounding the development of the next-generation Portals network programming interface and measurements of application sensitivity to network performance.

Accomplishments in FY15:

- Showed that offloaded networks have significant performance and energy gains versus on-loaded networks for power-efficient cores
- Provided SNL's Portals communication technology, which is the basis for new interconnection networks being designed and built by two hardware vendors for exascale systems

Planned Activities in FY16:

- Advance capabilities of Portals technology to support exascale-class systems
- Continue development of QThreads runtime capability to support asynchronous task parallel programming models

Resilience (SNL)

The next generation of computing platforms promises both new capabilities and increased capacity for meeting SNL's mission challenges. However, these platforms will involve new computer architectures. It is expected that the reliability of these systems may be degraded by both the sheer number of components as well as their susceptibility to errors as feature sizes are pushed to the limit. This project explores possible solutions to provide resilience to system errors that will enable our new ATDM codes to effectively use the new computational hardware.

Accomplishments in FY15:

- Demonstrated a technique that improves resilience to silent data corruption in iterative partial differential equation (PDE) solvers via self-correcting linear algebra operations in a parallel iterative PDE solver running on up to 256 cores
- Developed a node-aware data redundancy scheme for LFLR, a scalable resilience programing model

- Discovered and corrected significant scalability issues in the suggested next-generation resilience mechanism uncoordinated checkpoint/restart

Planned Activities in FY16:

- Complete an implementation of LFLR capabilities in an SNL ASC code for test and evaluation

Projects for the Input/Output, Storage Systems, and Networking Product (WBS 1.2.3.5.4)

The I/O, Storage Systems, and Networking 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, 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, 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 high-performance storage system (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 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.

Archive Storage (LLNL)

The Archival Storage project provides long-term, high-performance, archival storage services to ASC customers. This includes a collaborative software development effort (currently HPSS) between the tri-labs, ORNL, 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 includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and media, ongoing technology refresh, and data stewardship. HPSS provides scalable, parallel, archival storage interfaces and services to the tri-labs.[5]

A world-class array of hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms, thereby increasing computation. This includes disk arrays, robotic tape subsystems, servers, storage area networks (SANs), networks, and petabytes of tape media, enabling high-speed parallel transfers into a virtually unlimited data store.

Accomplishments in FY15:

- Continued ongoing HPSS software development and support, with focus on inspection and release of HPSS 7.5, which features partitioned metadata

[5] See http://www.hpss-collaboration.org/index.shtml.

- Began planning for production deployment of HPSS 7.5, including field performance evaluation to verify that existing HPSS core server platforms and metadata disk subsystems can exploit scalability of partitioned metadata without new hardware
- Developed and deployed hard quota and user notification features in archival quota system
- Began repack of nine-year old T10K Gen1 media (1-TB native) to T10K Gen2 media (8-TB native) to minimize data loss due to aging media and to reclaim slot capacity in libraries
- Deployed HPSS disk subsystems to increase disk cache capacity and disk file residency; deployed 40-Gbe network upgrades for increased archival performance
- Provided ongoing support for currently deployed archival storage systems, including selection, deployment, support, and maintenance of all archival storage hardware and media, customer and interface support, ongoing tech refresh, and data stewardship

Planned Activities in FY16:

- Continue ongoing HPSS software development and support, gathering requirements for design and development of a future major version of HPSS featuring support for multiple core servers
- Deploy HPSS 7.5 on existing HPSS core server platforms and metadata disk subsystems to exploit scalability of partitioned metadata
- Finish repack of T10K Gen1 media to T10K Gen2 media to minimize data loss due to aging media; eject/destroy Gen1 media to reclaim slot capacity in libraries
- Deploy secure computing facility (SCF) HPSS disk subsystems to increase disk cache capacity and disk file residency, and deploy 40Gbe network upgrades for increased archival performance
- Provide ongoing support for currently deployed archival storage systems, including selection, deployment, support, and maintenance of all archival storage hardware and media, customer and interface support, ongoing technology refresh, and data stewardship

Parallel and Network File Systems (LLNL)

The Parallel and Network File Systems (NFS) project provides for the development, testing (feature, capability, performance, and acceptance) and procurement of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included are the continuing development and support of Lustre as a fully featured file system for the range of ASC platforms, and the I/O support of various programming interfaces for parallel I/O.

This project develops and provides support for Lustre file system software. It actively works with the OpenSFS Lustre development community to add Lustre file system

scalability and reliability enhancements required by ASC platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

Accomplishments in FY15:

- Actively participated in and led the continuing community software development efforts of Lustre 2.5 and beyond
- Enhanced ZFS-based Lustre performance
- Actively participated in CORAL NRE file system and burst buffer working groups in preparation for the Sierra machine

Planned Activities in FY16:

- Support the development, testing, and deployment of new Lustre versions in classified and unclassified environments
- Enhance ZFS-based Lustre metadata performance in support of user and purge performance
- Update Lustre Monitoring Tool for use with Lustre 2.8
- Continue leadership of OpenSFS Technical Working Group
- Participate fully in CORAL NRE file system and burst buffer working groups in preparation for the Sierra machine

Networking and Test Beds (LLNL)

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.

Accomplishments in FY15:

- Tested Mellanox FDR IB with collective offload
- Acquired early access and conducted testing on RedHat ARM64
- Automated monitoring, detection, and correction of IB issues
- Conducted testing on Intel Haswell

Planned Activities in FY16:

- Integrate and test Mellanox EDR and Intel OMNI-Path with TOSS 3.0

- Work with RedHat on Power8 Little Endian bare hardware support
- Conduct hardware testing on OpenPower
- Conduct functionality testing of Intel Broadwell processor and Knights Landing
- Conduct functionality testing of Cavium and new Broadcom ARM64 hardware with RedHat 7

File Systems, Archival Storage, and Networking (LANL)

Capabilities of the Archival and File Systems components of the project include online file systems such as the NFS complex and enterprise-wide supercomputer file systems, global parallel file system (GPFS) development, deployment and management, scalable I/O (SIO) middleware development and support, interconnect technology development and deployment, SAN development and deployment, and archive.

The file systems element of the project provides end-to-end, high-performance networking and SIO infrastructure for the ASC program. 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. As the program moves toward exascale areas, these efforts will improve the scaling or programmability of the I/O in ASC applications for current and future large-scale machines. Current areas of investigation are scalable object stores, scalable indexing, burst buffer architectures, and scalable metadata.

Application Readiness capabilities are consolidated in this project, addressing issues with an application's production-run readiness 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 Application Readiness team identifies causes of unexpected behavior and deploys fixes in production so that system users are able to make productive use of the systems with their applications to solve their problems. The team provides production problem solving (create small problem reproducers, identify cause, consult with the relevant technical experts to find a solution, and verify the 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 and metrics, and analysis/profiling.

The project also includes software support capabilities focused on communication and networking libraries (MPI). The goal is to establish a strong development and analysis tool capability for current and next-generation HPC platforms, including parallel capabilities. The project 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.

Accomplishments in FY15:

- Identified the configuration changes that produced a 2X+ computational performance improvement running Pagosa on Sequoia

- Completed the functionality testing of high-performance I/O (HIO) on Trinity burst buffer test bed
- Designed a turquoise campaign storage solution using GPFS over ZFS (storage for files of short-term to mid-term lifetime)

Planned Activities in FY16:

- Research object store investigation, such as Scality, ViPR, and Ceph
- Continue HPSS, DB2 development for scalable HPSS metadata, RAIT development in HPSS, and parallel data transfer tool development for archive
- Continue emerging non-volatile memory research; investigate and assess promising devices, such as KOVE, HGST PCM, and the ULLtraDIMM
- Further file system I/O; support, test, and explore with middleware such as parallel log file system (PLFS) and multidimensional hashed indexed middleware
- Investigate resilience tools and analysis for tens of thousands of hard disk drives in a file system

Production Input/Output Services (SNL)

The Production I/O Services 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.

Accomplishments in FY15:

- Increased capacity of HPSS systems from 13 PB to 31+ PB on the Sandia Restricted Network (SRN) and the Sandia Classified Network (SCN) through technology conversion of tape cartridges from Oracle T10KA to T10KC versions
- Upgraded all servers for the Lynx data transfer services at SNL, including multiple servers on the SRN, Intersite HPC network, and classified network, which connects to ASC AT systems (new servers are running TOSS 2 and support the current Lustre file system clients)
- Contributed to the ASC tri-lab pre-exascale programming environment Level 2 milestone effort

Planned Activities in FY16:

- Upgrade HPSS servers to version 7.5.1 to provide enhanced performance in file creates and end-to-end data integrity
- Complete initial redesign effort for next-generation HPSS architecture (Version 8), preparing for exascale platforms

- Evaluate performance and reliability of 100 GEthernet link between SNL and LANL; evaluate data movement tools on the link and within the SNL advanced architecture test beds environment

Scalable Data Management (SNL)

The Scalable Data Management project provides critical R&D to support efficient and effective use of I/O capabilities on future exascale platforms. The research performed in this project addresses I/O concerns for application workflows, including efficient movement, management, and processing of data throughout the application workflow. This software will have tremendous impact on I/O for petascale and future systems because it allows for the creation of fully integrated scientific workflows that generate, analyze, and visualize data with minimal requirements for persistent storage. Current use of this technology includes data staging/caching to manage bursty I/O operations (for example, for checkpoints) and in-transit fragment detection for CTH, a co-design activity with the Scalable Data Analysis project.

Accomplishments in FY15:

- Used in-memory distributed hash tables as a way to communicate data in task-based programming models
- Actively investigated methods to develop better understanding of the I/O performance characteristics of SIERRA applications, both to provide guidance for current application deployments and to help prepare for migrations to future hardware and software platforms
- Played a key role in the joint ASC/ASCR workshops on storage systems and I/O; organized, participated, and ultimately authored large portions of the workshop report

Planned Activities in FY16:

- Complete and evaluate an I/O mini-application that represents SIERRA I/O patterns
- Deploy and evaluate Scalable I/O capabilities (Nessie/NNTI, Kelpie, Sirocco) on ASC platforms

Scalable File Systems (SNL)

The Scalable File Systems project provides R&D to investigate and develop file systems for future exascale platforms. Issues of particular importance are extreme scalability and resilience. To address these issues, this project is developing a file system that decentralizes management of devices to support a high degree of heterogeneity within a system of inherently unreliable networks and storage devices. The central components of this peer-to-peer-like system are “smart” servers that have access to a variety of different local and remote media (for example, disk, non-volatile random access memory (NVRAM), memory, and tape) and are pervasive throughout the computing platform.

These servers directly handle I/O requests, initiate third party transfers, or replicate the data as needed.

Accomplishments in FY15:

- Achieved significant performance improvement on throughput of the Sirocco networking infrastructure, addressing some of the critical concerns from the FY14 Level 2 milestone, in cooperation with SNL Production Computing Center
- Selected SNL's object storage solution as the common component moving forward for both Sirocco and Triton (the ANL file system project)

Planned Activities in FY16:

- Complete extended feature support for Sirocco, including the opportunistic gossiping protocols used to "piggyback" system information with normal data traffic, graph-based search capabilities within the connected storage components, and lock services
- Begin with ANL the implementation of the security model (originally designed as part of the ASC lightweight file system project)
- Explore/research strategies for leveraging heterogeneous storage tiers (for example, memory, NVRAM, disk) in Sirocco

Projects for the Post-Processing Environments Product (WBS 1.2.3.5.5)

The Post-Processing Environments 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, as well as 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, 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.

Scientific Visualization (LLNL)

The Scientific Visualization project conducts research and develops and supports tools for managing, visualizing, analyzing, and presenting scientific data. Research topics include topological analysis, particle visualization, and data compression techniques. Operational support for data analysis covers support of post-processing resources, including visualization servers, displays, 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. Display efforts include support of high-resolution, high-performance display devices for theaters and collaborative use areas. The project installs, maintains, and consults on software visualization tools, and supports demonstrations on the PowerWalls. The project maintains unclassified and classified 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, GPUs, and parallel storage systems. Hardware capabilities include three production visualization servers and several PowerWall clusters. A video display infrastructure drives PowerWalls and smaller displays. Visualization researchers continued to perform work in areas of topology, compression, and advanced data analysis techniques.

Accomplishments in FY15:

- Evaluated, selected, and procured visualization hardware to support ASC data analysis needs
- Completed upgrade of the Armadillo PowerWall theater with new projection systems and technology for seamless display
- Maintained the data analysis and visualization environment and provided operational support for all visualization facilities, including supporting projection equipment and facilitating the use of the data analysis clusters and associated storage
- Supported large-scale data analysis and visualization activities, including supporting ASC scientists with creation of visuals and movies for presenting and analyzing scientific data
- Exploited research results in data analysis and visualization for ASC simulations, including data compression, topological methods, and optimization of streamline tracing algorithm
- Developed a novel user interface for visualizing and querying data access patterns from streamline tracing application for both user-managed cache and system-managed memory map

Planned Activities in FY16:

- Continue to maintain the data analysis and visualization environment across LC platforms
- Provide operational support for all visualization theaters
- Upgrade the classified Tilden Room visualization theater with the same technology deployed this year in the large unclassified theater
- Support CCC, Grand Challenge, and other LC users with visualization and data analysis activities, including creation of visuals and movies for presenting and analyzing scientific data
- Pursue research activities in data compression, topological methods, and delivery of data analysis and visualization algorithms through the Lorenz MyLC portlet framework

Scientific Workflow and Data Management (LLNL)

The Scientific Workflow and Data Management project provides users with powerful and time-conserving ways to access, search, compare, and archive large-scale scientific data, and new high-level tools for managing the simulation workflow. This is achieved through the development of production-quality applications that enhance data management capabilities and the creation of innovative interfaces to job monitoring and vertical application frameworks.

Hopper and Chopper are the principal products of the data management effort. In the simulation workflow area, the Lorenz Web-based HPC application suite forms a foundation for providing new ASC-specific capabilities. Lorenz uses advanced Web technologies to make HPC more accessible, saving the user time while also helping the resources to be used more effectively.

Accomplishments in FY15:

- Released multiple new versions of Hopper and Chopper with a focus on incorporating MPI- and non-MPI-based parallelism for copy operations
- Investigated alternative fast metadata options for Lustre now that RobinHood has been ruled out
- Extended the MyLC dashboard to include more complete information about changes occurring within the center; provided users with a variety of ways to subscribe to this information
- Developed a suite of Lorenz portlets for viewing and analyzing Lustre-related I/O statistics by job, file system, and cluster

Planned Activities in FY16:

- Release new versions of Hopper and Chopper with a focus on automatically retrying failed copies, minimizing resource usage, and improving usability
- Incorporate an alternative, faster metadata option for Lustre file systems
- Extend the MyLC dashboard to support a focus-area view, where the user can display just items related to system status, Lustre performance, account information
- Implement Lorenz components for viewing data related to LC operations, for example, from the Center Monitoring project

Visualization and Data Analysis (LANL)

Data analysis and visualization are key capabilities in taming and understanding the ever-increasingly large datasets generated from extreme-scale scientific simulations. This project comprises research, development, deployment of software and facilities to production and ongoing expert support in this.

The production and facilities component of the project is to provide LANL weapons designers with visualization systems research and support. The project also provides individuals with expert knowledge in both visualization and weapons science to work directly with the LANL designers to utilize the full power of the hardware and software infrastructure for visualization and data analysis.

The project is responsible for both ParaView and EnSight visualization and data analysis software, including verifying the installations laboratory-wide and providing local user support in the use of the software. The project acts as a bridge between the LANL design community and the two vendors, Kitware and Computational Engineering International.

The project develops new visualization algorithms and systems to meet current and future capability requirements for ASC simulations. ASC simulations are currently producing massive amounts of data that threaten to outstrip the ability to visualize and analyze them. Therefore, it is important to develop and implement new techniques that enable working with these large datasets. Examples include in-situ analysis and visualization and database-driven (Hadoop/MapReduce based) post-processing analysis and visualization.

Accomplishments in FY15:

- Added initial ParaView Cinema capability that was successfully tested on TLCC-2 clusters with xRage
- Developed the initial prototype of ParaView Catalyst Adaptor for the LAP code

Planned Activities in FY16:

- Investigate automated techniques that detect events, measure their statistical significance, and capture relevant knowledge
- Continue application of in situ advances to ASC codes of interest
- Research advanced architecture of visualization systems, including incoming AT systems and software architecture for interactive users
- Support complex user visualizations and training of IC users with post-processing, in situ, and workflow automation for LANL production systems

Scalable Data Analysis (SNL)

The Scalable Data Analysis project provides data analysis capabilities and support for a range of SNL ASC customers—from analysts and code developers to algorithm designers and hardware architects. Capabilities include data manipulation, data transformation, and data visualization that contribute to insight from computational simulation results, experimental data, and/or other applicable data. A project emphasis is to deliver and support scalable capabilities that support increasing data sizes, data sources, and platform processor counts for ASC complex applications and system architecture.

This project includes production deployment and support services that enable ASC customers to carry out data analysis on ASC systems. This includes porting and installation of tools onto production systems; maintenance, testing, debugging, refinement and integration of tools in the end-to-end system environment as needed to assure effective end-user capabilities; and user support. SNL priorities include a focus on delivering and supporting analysis capability for Cielo and subsequent ACES platforms.

Current tools include scalable data analysis software released open source through ParaView and the Visualization Toolkit (VTK), an early-release in-situ data analysis library (Catalyst) for coupling directly with running simulation codes, and R&D prototypes for the analysis of results from ensembles of simulation runs. Current hardware platforms for data analysis are limited to data analysis/visualization partitions on the compute platforms with an emphasis on delivery of visualizations to desktop.

Partnering with ASC customers and other product areas, this project will continue to build on its successful ParaView and VTK-based products. The project performs R&D that advances these capabilities as needed for evolving next-generation architectures, ensuring that ASC's investment in data analysis and visualization will provide advanced capabilities on platforms from Cielo through future exascale systems.

Accomplishments in FY15:

- Developed in-depth memory tools to reduce in situ analysis overheads
- Developed ensemble analysis tools to support model tuning through comparison of simulation results with experimental data
- Led a tri-lab effort to identify analysis and visualization requirements for the tri-lab Pre-Exascale Environments Level 2 milestone

Planned Activities in FY16:

- Demonstrate and evaluate advanced analysis, visualization, and I/O capabilities for the SIERRA Toolkit to support the FY16 Level 2 milestone
- Prototype and evaluate integration of many-core visualization and analysis algorithms on key architectures, in particular the Intel KNL
- Continue ParaView and Catalyst releases, with production support

Projects for the Next-Generation Computing Technologies Product (WBS 1.2.3.5.6)

The Next-Generation Computing Technologies product includes costs for the planning, coordinating, and executing of the next-generation computing technology R&D activities. These activities will prepare the ASC applications and computing environment for the next computing paradigm shift to extreme parallelism, via heterogeneous and/or multi-core nodes.

Next-Generation Computing Environment (LLNL)

The next-generation computing enablement efforts will help prepare ASC for the post-petascale era, addressing the software environment and platforms. It includes gathering requirements and coordinating next-generation activities internally and externally. The efforts will enable application work to develop benchmarks for new commodity platforms as well as to adapt codes to the new platforms. The software efforts are focused on an environment that includes system-level software, resource management, development tools, data analysis tools, and programming models, while addressing ASC application requirements. Development environment efforts include work on debugging, power, and resilience. On the hardware side, efforts include tracking and collaborating on technology innovations. This effort includes interactions with vendors and academia, including planning and technical coordination for vendor contracts.

Accomplishments in FY15:

- Conducted co-design activities with vendors
- Executed LLNL portion of the tri-lab Level 2 milestone to provide feedback to vendors on key bottlenecks in performance for next-generation technologies
- Provided technical coordination for ASC AT system contracts
- Participated in standardization efforts for existing and evolving programming models, in particular the MPI forum (LLNL chaired the forum and Tools working group) and the OpenMP language committee
- Actively engaged with new efforts on programming models, including DOE ASCR X-Stack projects, PSAAP II center activities, and other academic efforts; coordinated with similar efforts at LANL and SNL

Planned Activities in FY16:

- Conduct testing of the next generation resource management Flux framework on CTS-1 systems at LLNL
- Continue active involvements in standardization efforts for MPI and OpenMP

- Continue and expand investigation of new programming models, especially task-based models, in close collaboration with the code team, next-generation efforts, and the AAPS team

Future Architecture Planning and System Requirements (LANL)

The major focus of the Future Architecture Planning and System Requirements project is to define requirements and potential system architectures for advanced systems platforms that meet ASC programmatic requirements and drivers. This project covers all aspects of program and procurement planning for current and advanced systems and strategic planning for supporting infrastructure. Additionally, this project provides a focus for the various planning efforts. In FY16, this project will focus on the project management of the ASC systems called Trinity and Crossroads. The focus in this project also includes the execution of DOE Order 413.3b.

Accomplishments in FY15:

- Provided program and project management for computing platforms, including requirements gathering and analysis
- Developed infrastructure plan design to support Trinity
- Participated in site-wide planning for power and cooling upgrades for future systems

Planned Activities in FY16:

- Continue to provide program and project management for computing platforms, including requirements gathering and analysis
- Develop infrastructure design in anticipation of CTS-2 and Crossroads systems
- Develop 10-year projections for power and cooling of projected LANL HPC systems for site-wide planning of utility power and water

Future-Generation Computing Technologies (LANL)

This project includes high-risk, high-reward research for future systems, including research on advanced programming models and hierarchical storage technology for advanced in-situ, and support of DOE initiatives such as FastForward and DesignForward.

The objective of the PINION project is to investigate the use of high-level data parallelism in the implementation of physics algorithms of interest to the ASC Program. With this model, algorithms are written using a relatively small set of data-parallel primitive operators, such as transform, reduce, and scan, along with custom functors. Backend implementations of the data-parallel primitives optimized for specific architectures then allow these higher level physics codes to be portable across these architectures, making efficient use of multi-core and many-core parallelism available on

each. The investigations undertaken by the PINION project are expected to have high relevance for the future of ASC codes.

Hierarchical storage technology for advanced analytics will investigate alternative hierarchical storage technologies, such as novel burst-buffer designs for support of in-situ and in-transit analysis and streaming analytics. Interactions with vendors for FastForward and DesignForward are included, such as gathering of computation and communication traces and support of mini-applications. Additionally, this project investigates virtual cluster technology that provides data-rich turn-key environments for specific applications.

Accomplishments in FY15:

- Generated polygons for 3D two-material interface reconstruction using data-parallel Pinion code, for each individual cell of a structured mesh
- Completed initial investigation of remote virtual cluster environments with open multiphysics application

Planned Activities in FY16:

- Investigate turn-key virtual cluster environments for HPC applications, which will be extended from previous work
- Continue development of a physics multimaterial AMR with PINION
- Investigate data streaming for applications on KNL and other architectures
- Use burst buffers in alternative design and novel configurations in support of in situ and in transit analytics with production codes

Architecture Office (SNL)

The objective of this project is to analyze potential computer and system architectures for platforms that meet future ASC programmatic requirements for ATS-3 and beyond. The primary activity is to establish a technology foundation for ASC to influence the directions for future hardware and system software architectures for ASC AT systems. The project will track HPC computer industry hardware/software trends with a specific focus on the identification of opportunities to influence future hardware architectures and development of future system software that provides an on-ramp for the ASC application code base. This project is also the focal point for the active collaboration of SNL technical staff with industry FastForward, DesignForward, and potential future ECI PathForward R&D projects.

Accomplishments in FY15:

- Hosted vendor visits at CSRI January 20–23 and February 17–19 (vendors included Nvidia, Intel, AMD, Cray, IBM, Broadcom, HP, ARM, and NanoPrecision); met with SGI and storage companies, including Intel, DDN, NetApp, SeaGate, Saratoga Speed, and Akonia Holographics, April 13–16 in Oakland, CA

- Supported development of CD-0 and the draft RFP for the Crossroads System Acquisition Project

Planned Activities in FY16:

- Complete CD-1, release Crossroads system RFP, support technical and business reviews, and complete vendor selection for initiation of the Crossroads system procurement contract
- Support current FastForward2 projects, DesignForward projects, and new ECI PathForward R&D efforts

Advanced Architecture Test Bed Research and Development (SNL)

This project will address a critical need for a range of experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy applications, enable application performance analysis with Mantevo proxy applications, support the Heterogeneous Computing and Programming Model R&D, the Software and Tools for Scalability and Performance projects, and for SST validation efforts. These test bed systems are made available for "test pilot" users who understand the experimental nature of these test beds. At the present time, it is more important to explore a diverse set of architectural alternatives than to push large scale. Discussions continue with Intel, AMD, IBM, NVIDIA, Micron Technology, and other computer companies regarding ASC interest in obtaining early access to experimental architecture test beds. These partnerships will establish a strong foundation for co-design activities that can influence future hardware designs.

Accomplishments in FY15:

- Received at SNL HP's first Moonshot system with production 64-bit ARM processors and presented analysis of tri-lab proxy applications with the pre-production version of these processors at ARM TechCon'14
- Demonstrated strong application performance on tri-lab mini-applications on SNL's POWER8 test bed

Planned Activities in FY16:

- Assess impact of advanced memory architectures on ASC codes
- Update advanced architecture test beds with technology refresh; publish report and user guide on capabilities for R&D users

Application Performance Analysis (SNL)

The purpose of the Application Performance Analysis project is to develop tools, techniques and methodologies to support the analysis and evaluation of current and next-

generation HPC technologies. A primary focus area of the project is to provide leadership for the Mantevo[6] project, and to facilitate the use of Mantevo mini and proxy applications as a tool. In addition, the project will utilize classic empirical and mathematical performance analysis methods to achieve its goals.

Next-generation computing platforms are expected to present significantly different architectural designs from previous platforms. In preparation for these changes, the project will explore the potential computing environments from processor to node, to inter-node. Mini and proxy applications, test beds, simulation capabilities provided by the SST, abstract machine models, and analytic performance models will be used. The outcome will be a better understanding of the characteristics and capabilities within the context of the computational science and engineering simulations of interest to the ASC program on emerging and future architectures and will inform hardware and software requirements. A primary activity for FY16 is to study and identify key performance issues of applications executing on emerging technologies and, in particular, the Trinity AT system.

Accomplishments in FY15:

- Characterized the key aspects of runtime performance for three mission-critical codes and presented findings at the Trinity CoE kick-off meeting at LANL
- Delivered a report on which benchmarks and proxy applications delivered to the broader DOE community effectively capture ASC code needs

Planned Activities in FY16:

- Demonstrate integration of performance improvements identified via mini-applications into application codes
- Support the Trinity and Sierra CoE through performance analysis and improvements for proxy and production applications

[6] https://software.sandia.gov/mantevo/

Appendix G: Facility Operations and User Support Subprogram (WBS 1.2.3.6)

This sub-program provides both necessary physical facility and operational support for reliable, cross-lab 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 local area network (LAN)/wide area network (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.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2014, and through quarter 3, fiscal year 2015, are reflected below for the FOUS subprogram.

- Provided ongoing support of unclassified and classified TLCC-2 and AT system compute platforms, parallel files systems, NFS, and infrastructure systems (LLNL)
- Completed deployment of new NAS home directory hardware in open computing facility (OCF) and SCF centers (LLNL)
- Implemented several monitoring improvements, including automated downing of IB links with high retry counts; eliminated cascading effects of bad links on the parallel file system, SAN, and attached platforms (LLNL)
- Completed upgrade to RSA authentication manager 8.1 on unclassified, SCF, and SNSI networks (LLNL)
- Completed migration of Linux-based security services to VM infrastructure (LLNL)
- Awarded construction contract for Building 654 modular unclassified computing facility; broke ground on May 28, 2015 (LLNL)
- Completed the Strategic Computing Complex (SCC) Computing Cooling Equipment project, providing highly efficient water-cooling capability for future compute platforms (LANL)
- Completed the facility power and cooling reconfiguration required for integration of Trinity (LANL)

- Completed the installation and integration of the next-generation backbone, improving network bandwidth between compute platforms, infrastructure servers, and file systems (LANL)
- Integrated and supported acceptance of the Trinity test systems Gadget and Trinitite (LANL)
- Refreshed and substantially improved the internal HPC documentation Web site to provide a simpler and more streamlined interface (LANL)
- Deployed new file systems in restricted and classified networks and retired older disk controller technologies to improve reliability and performance (SNL)
- Installed Trinity application readiness test bed Mutrino and assisted in monitoring and evaluation of liquid cooling design on this platform and Trinitite (SNL)
- Developed Trinity usage model document to describe the Trinity platforms computational environment and capabilities for tri-lab users (SNL)
- Placed 100 GE link between LANL and SNL into production state; began conducting tests with data transfer tools and measuring reliability (SNL)
- Demonstrated viability of bare metal provisioning and allocation of partial systems for Emulytics applications on Dark Nebula (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Further Automate Planned Cluster Maintenance to Minimize System Downtime during Maintenance Windows		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: A set of scripts will be written and deployed to further standardize cluster maintenance activities and minimize downtime during planned maintenance windows.		
Completion Criteria: When the scripts have been deployed and used during planned maintenance windows and a timing comparison is completed between the existing process and the new more automated process, this milestone is complete.		
Customer: LC		

Milestone (ID#TBD): Further Automate Planned Cluster Maintenance to Minimize System Downtime during Maintenance Windows
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.
Supporting Resources: LLNL CSSE and FOUS staff

Milestone (ID#TBD): Building 654 Beneficial Occupancy		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The construction process will be complete enough to allow the program to begin to install equipment necessary for siting CTS-1. Beneficial occupancy of the building will be attained.		
Completion Criteria: The program will be allowed to take beneficial occupancy of Building 654 after a project review of construction contract milestones.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: FOUS staff, Facilities and Infrastructure staff, vendor services contract, outside purchase order subcontracts		

Milestone (ID#TBD): Deploy a Common Computing Environment for Commodity Technology System-1 Platforms		
Level: 2	**Fiscal Year**: FY16	**DOE Area/Campaign**: ASC
Completion Date: 9/30/16		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: Successfully deploy TOSS on a CTS-1 production platform, providing a responsive and more efficient infrastructure to support computing for QMU and predictivity.		
Completion Criteria: Deployment of CCE on CTS-1 systems.		
Customer: ASC CTS-1 users		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: Tri-lab CSSE and FOUS staff		

Projects for the Collaborations Product (WBS 1.2.3.6.1)

The Collaborations product provides programmatic support for collaboration with external agencies on specific HPC projects. This product also includes collaborations with internal or external groups that enable the program to improve its planning and execution of its mission.

Program Support (LLNL)

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 tri-lab and LLNL-only activities, including collaborations with academic, industrial, and other government agencies.

Accomplishments in FY15:

- Continued FY14 procurement support, contract management, and next-generation computing planning
- Organized and hosted Principal Investigator (PI) meeting
- Supported annual HPC Best Practices meeting with Office of Science, now titled *HPC Operations Review Meeting*
- Supported bi-annual Predictive Science Panel (PSP) meetings
- Supported Presidential Early Career Award for Scientists and Engineers (PECASE) awardee
- Managed PSAAP II program (composed of six university centers chosen to collaborate/interact with NNSA laboratory personnel on R&D topics of interest to ASC)
- Supported SC15
- Supported FASTMath/ParaDiS; and SUstained Performance, Energy & Resilience (SUPER) Institutes
- Supported SCIDAC (Scientific Discovery through Advanced Computing) projects HAXTON and NUCLEI

Planned Activities in FY16:

- Continue FY15 procurement support, contract management, and program planning
- Support annual HPC Operations Review meeting with Office of Science
- Support bi-annual PSP meetings
- Provide technical project support of PSAAP II centers

- Support SC16; FASTMath/ParaDiS; and SUPER Institute
- Support SCIDAC projects HAXTON and NUCLEI

Program Support (LANL)

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.

Accomplishments in FY15:

- Supported PSAAP II collaboration efforts at LANL
- Published the ASC eNews newsletter
- Participated in PI and PSP meetings
- Provided LANL support for HQ

Planned Activities in FY16:

- Support PSAAP II collaboration efforts at LANL
- Publish the ASC eNews newsletter
- Participate in PI and PSP meetings
- Provide LANL support for HQ

Program Support (SNL)

The Program Support project provides critical coordination and integration activities essential to the success of ASC. It is divided into two distinct parts: 1) provide ASC programmatic planning, reviews, and communications; and 2) provide SNL outreach to the other institutions and programs.

This capability is critical to the ASC SNL program integration, communication, and management within the laboratory and with the external community. A significant management and integration function in this project is captured in the Leidos contract that provides support for NNSA HQ and SNL in communications and logistics. External advisory boards supported through this project also provide feedback to the ASC leadership team regarding the maturation of the predictive engineering sciences capability and the quality of SNL's computational science R&D. Support of external collaborations, including PSAAP II and the exascale initiative (with DOE/SC), is also included in this project.

Accomplishments in FY15:

- Organized and hosted sixth Predictive Engineering Science Panel (PESP) meeting and side meetings of PESP sub-panels

- Managed the Leidos contract to provide various administration support for HQ
- Supported ASC PI meeting in Monterey
- Organized attendance, booth, and meeting logistics for Supercomputing 2014 (SC14) Conference
- Supported programmatic needs of the PSAAP II program and the DOE Exascale Initiative

Planned Activities in FY16:

- Organize and host seventh PESP meeting
- Support external review panel meetings for Qualification Alternatives to the Sandia Pulsed Reactor (QASPR), the Engineering Sciences External Advisory Board, and the Computer and Information Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program and DOE Exascale Computing
- Support programmatic needs of NNSA tri-lab ASC Program and ASC executive committee, including programmatic meetings and communications management of the Leidos contract for program administrative support for HQ

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

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 provides the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources. Development and deployment of software to support the solution of manufacturing problems is also supported by the project. Visualization techniques that can be utilized in the Y-12 network and computing infrastructure will be evaluated and implemented. Finally, participation in NWC ASC-related activities is covered.

Accomplishments in FY15:

- Imported a complex Creo model of a proposed uranium process facility (UPF) operating area into Visionary Render and scripted to show common maintenance training operations; added a virtual iPad type interface that could be positioned within the simulation and used to access training videos and procedures.
- Developed a process to convert point cloud data generated by the Laser Detection and Ranging (LADAR) system currently available in Facilities Engineering, since 3D computer aided design (CAD) models are not available for the majority of facilities requiring training; converted laser scans of the New Hope History Center into

polygon meshed 3D models then integrated and imported into Visionary Render; tested a number of software components and importers to produce the best 3D models. Imported a 3D model of a shipping container into the laser-scanned environment to demonstrate how to integrate 3D CAD models and laser scanned images; scripted the container to show how an operator could sequentially remove and manipulate components of the container.

- Utilized COMSOL (a commercial software product) analysis to investigate safe shutdown technologies with potential intellectual property implications

Planned Activities in FY16:

- Relocate the Virtual Immersive Training Environment laboratory to a new location and upgrade the computing hardware to an HPC machine capable of handling the large point cloud data sets and 3D images.
- Complete a LADAR scan of a process area with enough detail for a training simulation environment. Convert the laser-scan images into polygon meshed model and import it into Visionary Render. Script the environment to show a current training scenario; demonstrate the virtual immersive training environment to include the integration of the optical tracking system and the head mounted display.
- Replace the aging classified simulation model terminal servers with new high-performance equipment. Set up the access and process for Pantex to run classified discrete event simulation models on the Y-12 terminal servers.
- Continue to collaborate with other sites (NWC colleagues) to best utilize codes and deploy capabilities as necessary to enable collaboration and to solve production manufacturing problems.
- Participate in National Security Enterprise ASC activities.

Projects for the System and Environment Administration and Operations Product (WBS 1.2.3.6.2)

System and Environment Administration and Operations product provides requirements planning, initial deployment, configuration management, and ongoing operational support for reliable production computing and storage environments. Activities include system and network administration and operations, user support, hardware maintenance, licenses, and common tri-lab computing environment integration and support.

System and Environment Administration and Operations (LLNL)

This project provides necessary operational support for reliable production computing environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support. Included within the scope of this product is the operational support for systems used as part of partnerships with academic, industrial, and other governmental agencies.

Accomplishments in FY15:

- Completed deployment of new NAS home directory hardware in OCF and SCF centers
- Provided ongoing support of unclassified and classified TLCC-2 and AT system compute platforms, parallel files systems, NFS, and infrastructure systems
- Implemented several monitoring improvements, including automated downing of IB links with high retry counts (this change eliminates cascading effects of bad links on the parallel file system, SAN, and attached platforms)
- Implemented improved file system monitoring via Splunk data analysis of system

Planned Activities in FY16:

- Integrate CTS-1 systems on classified and unclassified networks
- Continue to provide 24x7 support for TLCC-2 systems, Sequoia, and Vulcan
- Provide support for CTS-1 systems
- Retire Sierra, Juno, and Muir
- Select replacement identity management tool and begin conversion of existing workflows
- Using Splunk, develop the ability to correlate various authentication logs to identify issues and potential threats

- Migrate the remaining non-commercial RHEL-based security services to TOSS systems
- Deploy file system resources in support of CTS-1
- Deploy new non-"purgeable" project space in OCF and SCF environments

Hotlines and System Support (LLNL)

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC computing resources for the tri-lab as well as academic and industrial collaborations. This project includes computer center hotline and help desk services, account management, Web-based system documentation, system status information tools, user training, incident management systems, and application analyst support. Services are provided to both LLNL users as well as users from external sites, including LANL, SNL, and the ASC Alliance sites.

This project provides accounts administration, technical consulting, and documentation and training to facilitate the effective use of LLNL HPC systems. An accounts specialist team provides all account management services necessary for users to obtain accounts and access LLNL HPC systems. This includes account creation and removal, bank allocations, token management and visitor tracking for foreign national users. The technical consultant team provides technical support to LLNL users to enable their effective use of LLNL HPC systems. Consulting services vary from helping new users configure their environment, assisting experienced users with optimization of codes, and supporting other LC staff with monitoring of file systems, batch queues, and user environments. Extensive Web documentation, user manuals, technical bulletins, and training are provided to users via email, Web, and in-person training.

Accomplishments in FY15:

- Continued to provide ongoing support services for hotline operations, documentation, and training
- Continued the migration, reorganization, and consolidation of existing user documentation to the Confluence server
- Educated hotline staff members in accelerator technologies and their software libraries
- Began providing trusted agent support for National Security Systems Public Key Infrastructure (PKI) on the Enterprise Secure Network
- Became familiar with the Service Now Incident management software in preparation for migrating to it from Front Range in FY16

Planned Activities in FY16:

- Continue to provide ongoing support services for hotline operations, documentation, and training

- Finish the migration, reorganization, and consolidation of existing hotline tools and documentation to the Confluence Server
- Reorganize and update both internal and external LC Web sites
- Continue to develop in-house knowledge and expertise in accelerator technologies and their software libraries in preparation for providing support to the user community in their use
- Begin migration from the existing Front Range Incident tracking system to Service Now Incident tracking system

Facilities, Network, and Power (LLNL)

The Facilities, Network, and Power project provides for the necessary physical facilities, utilities, and power capabilities to ASC systems. Work in this area includes adequate raised floor space, flexible cooling solutions, and power to site large-scale ASC platforms. In addition, this project funds needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff. Also included are classified and unclassified facility networks, wide-area classified networks, and ongoing network operations. This project also enables enhanced collaborations with academic and industrial partners.

Accomplishments in FY15:

- Awarded Building 654 construction contract for a modular unclassified computing facility; broke ground on May 28, 2015
- Moved out of Building 115 and returned it to the institution
- Removed Dawn from computer room and restored floor to original condition
- Fixed unclassified server room floor and completed reconfiguration
- Installed new 20-ton chiller in preparation for Sierra

Planned Activities in FY16:

- Evaluate the reuse of LLNL groundwater for tower water fill; evaluate the use of water conservation techniques to reduce blow down at cooling towers at Building 453
- Complete demolition of under floor for CTS-1 system siting
- Complete power installation for CTS-1 system
- Complete Sierra facility preparation projects: Building 453 Power Modernization and Advanced Development Required for Cooling Loop

System Administration and Storage (LANL)

The System Administration and Storage project covers all services for computational 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 include system configuration, system and user security, resource management, system administration and monitoring, archival storage, parallel storage, and NFS.

Accomplishments in FY15:

- Supported HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers
- Initiated efforts to develop a high-capacity, short-term data archival capability, providing a mid-tier storage capability to fill the gap between high-speed parallel file access and low-speed tape archival
- Ensured workloads were properly executed on compute resources, through configuration of queues and scheduling policies, daily monitoring, and problem resolution
- Specified and procured file transfer servers, network switches, NFS filers, and other equipment, to replace end-of-life systems
- Completed upgrade of systems' network capabilities in concert with the next-generation backbone (NGBB), to meet the performance requirements of the Trinity platform

Planned Activities in FY16:

- Support HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers
- Provide administration and resource management support for Trinity-Haswell during the open science campaign, the transition to the secure network, and the first secure ATCC
- Provide configuration and administration support for the CTS-1 platform
- Provide continuing security enhancements and requisite software stack upgrades on deployed compute platforms and infrastructure servers

Operations and Procurement Support (LANL)

The Operations and Procurement Support project provides around-the-clock operations and monitoring of the scientific computing resources, including performance computers such as Trinity, Luna, ViewMaster II, Cielo, Moonlight, and data storage and retrieval

systems such as HPSS. In addition to monitoring all components 24x7x365, the computer operators provide systems hardware maintenance for all ASC platforms. Working with the vendor system engineers, the operators also provide backup hardware support for the Cielo capability system. This includes all components of the production computing environment, from compute engines, hardware, file servers, archival storage systems, the facilities they reside in and utilities they are dependent upon, to all required software on these systems.

The procurement support aspect of this project assists customers with the technical and administrative aspects of planning, procurement, and contractual agreements for computer hardware and software products and services.

Accomplishments in FY15:

- Provided 24x7x365 first-level, on-call support, operations, and monitoring of compute and file system resources
- Provided continuing hardware maintenance, assessment of system life spans, and guidance on hardware maintenance burdens for computing equipment
- Supported installation of the Trinity-Haswell platform, Trinity test systems, and Trinity file systems
- Supported the acquisition of several mission-essential equipment components

Planned Activities in FY16:

- Provide 24x7x365 first-level, on-call support, operations, and monitoring of compute and file system resources
- Provide continuing hardware maintenance, assessment of system life spans, and guidance on hardware maintenance burdens for computing equipment
- Fold the Trinity platform into the operations, monitoring, and maintenance program
- Support installation of Trinity-KNL platform
- Support installation of the CTS-1 platform

Computing Platform Integration and Deployment (LANL)

The scope of the Computing Platform Integration and Deployment project is to accept delivery and begin deployment of production CT system and AT system platforms. This includes participating in developing the design requirements as part of a tri-lab requirements planning team. Primary capabilities include completing the acceptance tests, diagnostics test, integrating the systems into the LANL unclassified network, system stabilization, and transition into the classified network. Included in this project is support for the ASC CT system acquisition strategy and provision for requirements that help to achieve the strategy.

The objective of the project is the integration of all hardware and software components to deliver a system environment to application users for programmatic work. This includes site preparation to prepare the SCC for deploying these production capacity systems. The integration and deployment activities will focus on the following areas: System/OS, File Systems, Interconnect, External Network including NGBB, Regression Testing, and Monitoring.

Accomplishments in FY15:

- Designed and executed a performance assessment of Cray's Sonexion file system, for acceptance of the Trinity platform
- Executed rigorous testing and evaluation processes of the Trinity test systems, in support of system acceptance
- Participated in the tri-lab CTS-1 proposal review and selection process
- Provided continuing integration and top-tier diagnostic and troubleshooting support for standing production compute platforms and file systems

Planned Activities in FY16:

- Integrate Trinity-Haswell platform into LANL's production computing environment
- Provide early integration and acceptance support for Trinity-KNL platform
- Transition the Trinity platform into LANL's secure, production computing environment
- Integrate CTS-1 platform into LANL's secure, production computing environment
- Provide continuing integration and top-tier diagnostic and troubleshooting support for standing production compute platforms and file systems

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

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 the ASC tri-lab, 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.

Accomplishments in FY15:

- Surveyed user community and successfully acquired actionable and constructive feedback on HPC resources, policies, and services

- Refurbished internal documentation Web site to streamline information delivery and simplify site navigation
- Negotiated a restructuring of parallel tools licenses to minimize program costs and better facilitate common usage patterns
- Completed and published Trinity system documentation

Planned Activities in FY16:

- Prepare CTS-1 system documentation and training materials
- Provide user support to early Trinity system users and to *Trinity Open Science Campaign* research staff
- Facilitate the transition of user data and utilization from legacy parallel storage systems to newly integrated Lustre file systems
- Provide ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure

Facilities, Networking, and Power (LANL)

The Facilities, Networking, and Power project is responsible for the engineering, design, operation, and maintenance of the mission-important electrical, mechanical, cooling, network services, and other computing infrastructure in support of the ASC program. The project provides support for infrastructure design upgrades, project and space management, user interface and oversight, demolition and decommissioning of older systems, network backbones, user LANs, classified/unclassified network hardware and services, DisCom WAN, and computer site preparation for new platforms. Because the tri-lab community requires the systems to be operational at all times, the project provides on-call support after hours and on weekends for facility related issues.

Accomplishments in FY15:

- Completed the SCC computer cooling equipment project
- Completed the physical infrastructure and integration for Trinity effort, allowing for warm-water cooling of the Trinity system; completed similar installations for cooling and infrastructure testing relating to the Trinity test systems
- Provided ongoing operations, maintenance, and configuration of electrical and mechanical systems for ASC computing platforms and facilities
- Decommissioned the legacy PaScalBB network backbone and fully integrated the replacement NGBB, providing latest generation and expandable networking infrastructure to accommodate future compute platforms

Planned Activities in FY16:

- Initiate the requirements assessment for, and design of, future power and cooling equipment, as required to accommodate next-generation compute platforms

- Provide power and cooling infrastructure for integration of the CTS-1 platform
- Provide ongoing operations, maintenance, and configuration of electrical and mechanical systems for ASC computing platforms and facilities

Production Computing Services (SNL)

The Production Computing Services project's goals are to operate and maintain all ASC production platforms and associated support systems, and operate 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 CT capacity systems into SNL's production computing environment, in collaboration with WBS 1.5.5.6 CCE. Support of CCE common service and environment decisions and configuration management activities are also provided.

This project has expertise in operating capacity computing clusters; integrating file servers at the system or facility-wide level; deploying new computing, storage, and data management platforms; and in retiring end-of-life platforms. System administration for complex HPC environments is provided, as are design and development activities for new innovative advanced architecture computing platforms.

Accomplishments in FY15:

- Deployed new file systems in restricted and classified networks; retired older disk controller technologies to improve reliability and performance
- Installed and deployed new test bed systems for general tri-lab and collaboration use (White, Shannon, Shepard, Moonshot HP system), incorporating Power 8, Power 9, ARM, and various accelerator nodes
- Installed Trinity application readiness test bed Mutrino and assisted in monitoring and evaluation of liquid cooling design on this platform and Trinitite
- Procured and installed new Arista high-speed switches to replace aged HPC network backbone in Building 880 Computing Annex

Planned Activities in FY16:

- Receive, integrate, and deploy into production two CTS-1 platforms
- Continue operations support of all production platforms and test bed systems

User Support (SNL)

The User Support project provides user support and associated resources for SNL computing systems and tri-lab resources. User support activities focus on improving the

productivity of the entire user community, local and remote, in utilizing the ASC HPC resources.

This project deploys and maintains the following SNL capabilities for user support: 1) coordination between user support activities and leadership in adopting Information Technology Infrastructure Library (ITIL) principles and practices; 2) ITIL-compliant incident, problem, and knowledge management tool set; 3) training facilities and equipment; and 4) a Web portal for HPC-related information, real-time data, and documentation.

In addition, this project provides the following user support capabilities in conjunction with other projects: 1) a tiered user support structure (HPC service desk) that responds to SNL and tri-lab user requests received via phone, email, Web-based requests, and in-person visits; 2) the Synchronized Account Request Automated Process (SARAPE) tri-lab account provisioning Web-based tool; 3) Web-based, classroom, and one-on-one training; and 4) direct support in utilizing ASC resources.

This project also funds the SNL user support team's involvement in collaborative efforts such as PSAAP II and ACES.

Accomplishments in FY15:

- Provided user support for SNL and tri-lab ASC computing, including Tier 1 issue response and application support such as run-time analysis, debugging, tuning and porting
- Continued partnering with LANL to provide user support for ACES platforms; deployed revamped user support in preparation for Trinity, including development of Trinity Usage Model
- Deployed a new change management process for use by all SNL HPC and HPC-related teams
- Continued to develop, maintain, and improve user support resources, including training facilities, user support and collaborative tools, and Web portals for HPC and ACES information

Planned Activities in FY16:

- Provide user support for SNL and tri-lab ASC computing systems
- Continue to develop expertise in support of next-generation architectures and software environments
- Finalize deployment of new HPC and ACES Web portals
- Continue ramp up to deliver user support for Trinity

Facilities, Networking, and Power (SNL)

The Facilities, Networking, and Power project funds the power and space charges assigned to HPC systems (capacity and file system servers) and long-term hierarchical storage servers (running the HPSS software product). It provides for facilities and personnel to manage installation and removal of computing platforms, file systems, visualization systems, networking equipment, power distribution systems, and cooling systems in support of all computing resources. It also funds major operations contracts such as the ASC Distance Computing (DisCom) WAN.

Facilities professionals have reduced overall operating expenses by minimizing cooling and electrical distribution expenses over the last several years through a comprehensive program of introducing more efficient computer room air conditioning units, using higher voltage electrical source power distribution units, exploring alternative energy sources and conservation mechanisms, which include reducing the volume of chilled water required for cooling and improving air flow in the facility by minimizing obstructions underneath the computer floor. These efforts have been recognized with several SNL-specific and national awards, including three 2011 EStar Awards from the DOE Office of Sustainability Support.

Accomplishments in FY15:

- Placed 100-GE link between LANL and SNL into production state; continued conducting tests with data transfer tools and measuring reliability
- Implemented evaluation effort to characterize direct operations cost comparisons between identical TLCC-2 systems, one fully air cooled and one partially liquid cooled
- Retired several older clusters to begin freeing up space and power for CTS-1 installations

Planned Activities in FY16:

- Prepare both the 880 Computing Annex and Building 725 for installation of CTS-1 platforms
- Renegotiate a long-term contract for DisCom service
- Engage in technology discussions with potential vendors of high-speed encryptors

Projects for the Common Computing Environment Product (WBS 1.2.3.6.3)

The goal of the CCE product is to enable a common environment across the tri-labs that was initially 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, including but not be limited to, OS software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues such as configuration management, licenses, WAN access, and multi-realm security.

System Software Deployment for Commodity Technology Systems

The projects involved in this area include the TOSS and monitoring/metrics integration.

TOSS is the software stack that runs on 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 OS capable of running MPI jobs at scale on CT system hardware. The system must meet CCE requirements for providing a common software environment on CT system hardware across the tri-lab complex, now and into the future. TOSS provides a complete product with full lifecycle support. Well-defined processes for release management, packaging, quality assurance 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.

The Monitoring and Metrics Integration project targets efficient and productive use of HPC systems as well as informed future planning through: 1) effective monitoring of all measurable or reportable conditions on compute platforms, both current and future, that can impact the performance of both applications and throughput on those platforms; and 2) appropriate transformation of monitored information into metrics and transport of those metrics to facilitate their use by system utilities, applications, resource managers, users, system administrators, and management. Integration of information from disparate data sources will enable greater system understanding and response to system conditions.

Accomplishments in FY15:

- As previously reported in the project "System Software Environment for Scalable Systems," released updates to TOSS (version TOSS 2.2.1–1, 2.2.1–3, and 2.3–3) that included Lustre version 2.5.3, security updates, and bug fixes; released TOSS 2.3 (based on RHEL 6.6, the latest release from Red Hat); developed TOSS 3 (based on RHEL 7)
- Developed/deployed a tri-lab-coordinated software engineering and release process, accommodating 1) OS development, integration, test, and distribution, from a central,

shared build farm; and 2) patch, module, and software package utilization of the central resource

- Developed testing infrastructure improvements (Pavilion)
- Tested trial deployments of the LDMS at the tri-labs
- Developed LDMS Version 2.4, with new and enhanced capabilities

Planned Activities in FY16:

- As proposed in the project "System Software Environment for Scalable Systems," provide ongoing TOSS software development and support; develop/deploy TOSS 2.X for legacy systems (based on RHEL 6.X); develop/deploy TOSS 3.X for CTS-1 systems (based on RHEL 7.X)
- Continue SLURM support efforts through tri-lab collaboration; begin investigation of running SLURM-only on CT systems
- Develop identified collaborative system software tasks, including investigation of new architectures (for example, GPGPUs and ARM), integration of virtualization, implementation of logging/monitoring improvements, and testing infrastructure improvements (Pavilion)
- Continue to improve the HPC monitoring, analysis, and feedback infrastructure within the tri-lab through collaborative development, enhancement, and deployment of relevant tools and practices

Programming Environment Development/Support for Tri-Lab Systems

The goals of the Programming Environment Project are to enhance productivity of the tri-lab application development teams, operation teams, and analysts by developing and deploying user tools and programming environments to support a variety of applications running on tri-lab HPC resources. Challenges include supporting changing processor and systems technology and evolving programming models for advance architectures. This project entails software integration, feature enhancements, installation, training, support for vendor provided tools, open source software tools, and lab-developed tools.

Focus areas include Performance Analysis Frameworks and tools that include the Component-Based Tool Framework (CBTF) and Open|SpeedShop (O|SS) that is built on top of it. Debuggers include scalability work with Rogue Wave to further scale and enhance TotalView, the Stack Analysis Tool (STAT), and other subset debugging efforts. The MPI integration and scaling efforts provide development support to the communities to add and fix features in both MVAPICH and Open MPI. Lastly, the Open Source Contract Maintenance effort provides funding to outside developers who maintain tools and tool infrastructures that are critical for code teams or serve as the basis for internal tools. Each contract includes support for all three laboratories, and all three laboratories in close collaboration provide the technical guidance for the three contracts. This currently includes tools such as Open|SpeedShop, TAU, MUST/Vampir, and Valgrind.

Accomplishments in FY15:

- Assessed current ASC support tools ecosystem in preparation for upcoming CT system and AT system architectures; identified gaps, initiated transition efforts, continued involvement in programming model and runtime R&D efforts; began initial exploratory efforts regarding ecosystem into ARM, including Open|SpeedShop and Dyninst components; began efforts with Tau and MUST
- Deployed new interface enhancements into the CBTF (through O|SS) to support broad-based usability and integration; deployed further memory and threading experiments to user communities
- Delivered enhanced debugging capabilities through the TotalView BIGCAT scalability effort with scoping completed for SOW III to better support Intel KNLs
- Provided MVAPICH and Open MPI enhancements and bug fixes, support for MPI_THREAD_MULTIPLE and native PMIx support; targeted completion of libfabric Open MPI project to deliver high performance Open MPI on Trinity

Planned Activities in FY16:

- Continue development and support efforts for debuggers, performance analysis tools, and MPI as programming models and architectures evolve
- Increase focus on updating functionality of the tool ecosystem to support upcoming AT system architecture needs, including accelerator support, port to OpenPower, Nvidia GPUs, and Intel's KNL
- Further develop capabilities and collaboration on the Spack HPC software package manager to enhance efficiency for support and developer environments
- Continue to study MPI performance tuning and further develop a STUdy Driver (STUD) system and test set to represent and routinely test MPI usage patterns
- Track and assess compiler development across various emerging architecture in support of advanced architecture software environment needs

High Performance Computing Environment Integration for Tri-Lab Systems

The HPC Environment Integration project targets the ability to work across sites with minimal transition and access restrictions. Differences in tri-lab security implementation and network restrictions as well resource access and authorization processes have been a hurdle. Efforts target network access infrastructure, cross-realm authentication and resource management and environment standardization. Current efforts include establishing a cross-site authentication and resource approval through enhancements to the SARAPE system. This is a Web-based application that allows users within restricted domains to request selected CCE resources to which they are permitted access. It addresses the APIs required to help interface SARAPE with other tools required to manage accounts among the tri-labs. As part of the inter-site HPC deployment, a service catalog will be deployed through which collaborators can view and request accounts and

services available in the shared environment. The Shared Workspace effort is the infrastructure for promoting collaboration across the laboratories. It currently includes the Gforge server that is housed and managed at SNL.

Accomplishments in FY15:

- Added reporting capabilities to WC Tool for new PSAAP II universities
- Delivered quarterly tri-lab workload characterization reports to NNSA (WC Tool)
- Released WC Tool version 2 at SNL
- Increased SARAPE process automation for SNL resources, including auto-validation and immediate transfer of clearance information from SARAPE to corporate database
- Incorporated Identity Provider (IDP) credentialing for expanded methods of authentication to SARAPE by tri-lab users, including recognition of HSPD-12 personal identity identification (PIV) from LLNL and SNL

Planned Activities in FY16:

- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments; investigate and develop streamlined ASC HQ reporting tools
- Continue to operate and improve the tri-lab SARAPE process for all remote access account requests and implement evolutionary improvements
- Provide enhanced MySARAPE capability at all three laboratories, to allow users access to account status at all sites
- Develop capability for tri-lab SARAPE host processing agents to obtain metrics from SARAPE

Monitoring and Metrics Integration for Tri-Lab Systems

For FY16, the Monitoring and Metrics Integration project was merged into the System Software Deployment project.

Accomplishments in FY15:

- Continued development of site specific Splunk filters
- Progressed in continuous monitoring interactions surrounding Pavilion and Scummee/LDMS interaction
- Tested and deployed initial trials of LDMS at LLNL, LANL, and SNL
- Deployed proof of concept of planned Trinity monitoring infrastructure (including use of SNL's Power API) on ART systems at LANL and SNL
- Released LDMS Version 2.4 with new and enhanced capabilities, which was included in April 2015 TOSS release

Planned Activities in FY16:

None.

Projects for the Special Purpose Facilities, Systems, Operations, and Support Product (WBS 1.2.3.5.4)

The Special Purpose Facilities, Systems, Operations, and Support product provides special purpose HPC resources to the DOE community and the necessary support and maintenance of these systems and facilities. This includes special security controls and special purpose facilities in addition to the standard HPC operations and support activities necessary to support these resources.

Special Purpose Computing (LLNL)

The Special Purpose Computing project at LLNL leverages the established expertise, resources, and practices of the ASC Program to provide robust computing services and software capabilities to specially tasked research and assessment personnel. The project seeks to optimize the utilization and performance of HPC resources within the particular security and capability requirements of the user community, to facilitate the transfer of latest generation technology into these unique computing environments, and to coordinate the integration and support of ASC-developed software tools and resources, as necessitated by user activities.

Accomplishments in FY15:

- Provided HPC procurement, system administration, and operational support
- Identified new HPC computer room space and created an MOU governing support of the same

Planned Activities in FY16:

- Provide HPC procurement, system administration, and operational support
- Commission new computer room space in Building 453

Special Purpose Computing (LANL)

The Special Purpose Computing project at LANL leverages the established expertise, resources, and practices of the ASC Program to provide robust computing services and software capabilities to specially tasked research and assessment personnel. The project seeks to optimize the utilization and performance of HPC resources within the particular security and capability requirements of the user community, to facilitate the transfer of latest generation technology into these unique computing environments, and to coordinate the integration and support of ASC-developed software tools and resources, as necessitated by user activities.

Accomplishments in FY15:

- Provided ongoing operations, maintenance, and administration of compute facilities, platforms, and file systems
- Provided ongoing support and expanded deployment of software resources

Planned Activities in FY16:

- Provide continued operation, maintenance, and administration of compute facilities, platforms, and file systems
- Provide ongoing support and expanded deployment of software resources

Special Purpose Computing (SNL)

The National Security Computing Center (NSCC) at SNL provides capability class computing, Emulytics and Data Analytics platforms, high performance file systems and long distance network access for customers engaged in special purpose projects residing in a high security environment. These services and platforms derive from products developed and deployed through the ASC program.

Accomplishments in FY15:

- Operated storage systems, archive systems, and production systems supporting the LLNL, LANL, and SNL NSCC programs
- Implemented hierarchical data storage system capability using StorageTek tape subsystem, 150 TB disc cache, and HPSS software
- Deployed encrypted file system service and capability to serve multiple communities of interest; developed security plan and usage documentation

Planned Activities in FY16:

- Retire Cielo del Sur system
- Continue NSCC operations
- Continue user support development for NSCC

Appendix H: Academic Alliance Centers

Accomplishments listed below represent those of FY15.

University of Utah

The Uncertainty Quantification-Predictive Multidisciplinary Simulation Center for High Efficiency Electric Power Generation with Carbon Capture

Accomplishments in FY15:

- Integrated full system computer science components and physics models across all levels of the validation hierarchy; completed a UQ analysis cycle for the full hierarchy
- Began first round of predictions on a new design for a 500-MW oxy-coal advanced ultra-supercritical coal boiler
- Released to open source Utah's runtime framework, Uintah 1.6.0, which now has been demonstrated to run and scale on all major DOE and NSF platforms
- Demonstrated two new visualization strategies that have promise for in-situ analysis and visualization of the ever-growing datasets
- Developed particle reaction, heat transfer, and transport subgrid models; implemented them to relax to the time and length scale of the large eddy simulation, allowing the multiphysics, particle-laden, radiating reacting-flow, applications code to robustly run at the large eddy simulation time scale

Planned Activities in FY16:

- Work with Alstom Power (the Utah center's industrial partner) to focus the validation hierarchy on short-term, critical design and operation questions
- Increase computational efficiency for production code and more efficiently utilize NNSA computing resources for the overarching problem
- Integrate in-situ visualization and data management tools into production code
- Revise development hierarchy to emphasize coal-ash/mineral-matter transformation process to quantify the wall deposit properties more fully
- Teach a one-semester course simultaneously at three universities (University of Utah, University of California at Berkeley, and Brigham Young University) and include interested national lab participants

University of Illinois, Urbana-Champaign

Center for Exascale Simulation of Plasma-Coupled Combustion

Accomplishments in FY15:

- Built a center with 20 PhD students, research staff, and faculty, all engaged in center-level goals and interacting with NNSA labs
- Developed an integrated simulation model that represents all known physical mechanisms for the plasma initiated and mediated combustion application, with initial cost of only twice a corresponding inert-flow turbulence calculation
- Constructed the target experimental application in an Illinois laboratory, with integrated turbulent flow, combustion, and two plasmas; measured it with multiple advanced diagnostic tools
- Applied a suite of CS tools to the principal prediction simulation platform PlasComCM: vectorization analysis with Vector Seeker, Just-in-Time compilation for runtime adaptation, overdecomposition with AMPI, MxPA multihardware (GPU) environment, and an auto-tuning application XPP that includes a prototype of a general annotation approach designed to facilitate tool interoperability
- Made a year-1 target prediction with the full model on the full application, with primary and secondary quantities of interests; identified key pacing uncertainties for going forward

Planned Activities in FY16:

- Demonstrate preliminary CS tool interoperability within annotation framework on full application simulation PlasComCM
- Increase physical fidelity of pacing uncertainties based on uncertainty analysis of current integrated physical model
- Develop secondary quantities of interest to evaluate fidelity for important physical interactions, both in physics-targeted bench-top experiments and in the full application
- Make a refined prediction with full application tool of the target application, which also advances combustion engineering through the use of plasmas
- Continue to foster close interactions with NNSA labs through student internships, staff and faculty visits, and the planned Workshop on Exascale Software Technologies

Stanford University

Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment

Accomplishments in FY15:

- Developed full exascale software stack, including Terra/Lua/Liszt/Legion integration and demonstration of Soleil-X software on CPU/GPU single-node
- Simulated fully developed turbulent duct with particles and radiation forcing on large-scale clusters; demonstrated scalability using Soleil-MPI software infrastructure
- Analyzed the effect of variable-size particles on the heat transfer to the fluid and quantification of the effect of measurement uncertainty on the tail of the fluid/particle temperature probability distribution function at the channel exit
- Developed verification solutions (analytical and manufactured) for single-particle transport under various flow and thermal radiation conditions
- Qualified the experimental apparatus in terms of both flow conditioning and particle injection characterization; carried out preliminary comparisons with Soleil-MPI predictions

Planned Activities in FY16:

- Demonstrate the exascale software stack and Soleil-X on a multinode system, involving multiple CPU and GPUs
- Simulate fully developed turbulent duct with particles and radiation forcing on large-scale clusters; demonstrate scalability using Soleil-MPI software infrastructure
- Develop and validate improved point-particle and radiation models to include the effect of accurate drag law estimation and light scatter respectively
- Demonstrate multilevel Monte Carlo algorithms for quantifying the uncertainty in particle-laden turbulent flows in the presence of a large number of uncertainties
- Carry out first round of validation of particle-laden turbulent flow in in-house experimental apparatus using Soleil-MPI; complete qualification of thermal radiation source in the in-house apparatus and demonstrate preliminary measurements of particle and fluid temperature at the exit

University of Florida

Center for Compressible Multiphase Turbulence

Accomplishments in FY15:

- Performed multiple 3D simulations of the demonstration problem (cylindrical explosive charge) for sensitivity analysis and UQ to detonation modeling; carried out a hero run that included simple models of all the relevant physics

- Established a hierarchical validation plan for the microscale, mesoscale, and macroscale; performed groups of simulations at the mesoscale that quantified uncertainties and revealed limitations of the 1D simulation code; developed a multifidelity surrogate framework and extrapolation techniques that together will facilitate UQ for both the science and exascale teams
- Built CMT-nek as a trunk of nek-5000, which will become the Center's main physics code as it progresses towards exascale computing; carried out 3D simulations of a sudden impulse over a random packing of spheres, which in the limit of an isolated sphere provides rigorous verification
- Focused on the development of foundational behavioral emulation methodologies and methods (coarse-grained multidimensional computation modeling), the design and implementation of lamport clock based behavioral emulation parallel discrete-event simulator prototypes in shared memory and HW (single-FPGA), and the validation of developed behavioral emulation methodologies of several many-core mesh devices with the parallel discrete-event simulator prototypes; developed abstract behavioral emulation application models for the CMT-nek spectral element solver for use in architectural and algorithmic design space exploration
- Developed a platform-independent code optimization framework based on auto tuning for the derivative computing kernel of CMT-nek; developed auto tuning methods for performance and energy optimization; conducted extensive benchmarking on a variety of architectures

Planned Activities in FY16:

- Bring the physics for the demonstration problem to the current state-of-the-art; improve the productivity of Rocflu for large scale simulations (towards 100k cores)
- Perform a detailed uncertainty budget for the demonstration problem; carry out other subscale validations and extend the multifidelity surrogate capability; carry out the collaboration with the exascale team
- Focus on extension of behavioral emulation methodologies beyond the device level to the node and system levels, the integration of network congestion models and UQ theory, the use of select existing coarse-grained parallel discrete-event simulator such as SST for behavioral emulation simulations, and the extension of parallel discrete-event simulator prototypes to provide increased simulation speed and scalability
- Investigate bottlenecks in the performance of CMT-nek when the optimized kernel is run as part of the main code; delve into thermal analysis of the code kernels and load balancing algorithms; incorporate thermal models and DVFS along with auto tuning methods and extend them for multicore machines; develop load balancing algorithms for PIC problems for hybrid multicore architectures
- Investigate sequences of fully resolved microscale simulations consisting of a random packing of particles subjected to more realistic time-dependent flows that progressively better approximate the actual problem of explosive dispersal; further investigate the effects of Mach number, volume fraction (using random particle

packing), and Reynolds number using a series of simulations; process the results to extract relevant information (peak forces, temperatures, velocity fluctuations) for model development

Texas A&M University

Center for Exascale Radiation Transport

Accomplishments in FY15:

- Demonstrated that the radiation transport code, PDT, could scale to 393,216 processors with 70 percent efficiency; developed an optimal scheduler for PDT and improved a parallel performance model for PDT, which played a major role in this achievement
- Implemented basic nested parallel capability in the general-purpose parallel services library, STAPL
- Completed initial implementation of LLNL transport mini-application, KRIPKE, with the STAPL library
- Obtained detailed agreement between PDT and the community-standard continuous-energy Monte Carlo transport code, MCNP, after substantial work on the generation and verification of multigroup cross-sections for PDT simulations (including discovery of a bug in the NJOY cross-section processing code)

Planned Activities in FY16:

- Expand and refine the capabilities for nested parallelism in STAPL
- Employ nested parallelism in PDT and determine its efficacy for radiation transport calculations
- Begin the extension of the optimal parallel sweep algorithms to unstructured meshes for radiation transport calculations
- Install, test, and characterize the neutron generator and data acquisition system; begin using them in experiments
- Complete experiments and simulations for calibration of the Impurity Model-1; complete validation study for Impurity Model-1

University of Notre Dame

Center for Shock Wave Processing of Advanced Reactive Materials

Accomplishments in FY15:

- Developed transient PGFem3D solver and quasi-steady-computational theory of homogenization

- Developed Parallel Asynchronous Space-Time Algorithm using a Domain Decomposition Method (PASTA-DDM) with non-matching grids
- Developed stable HPX RTS; released to the public HPX-5 version 1.0
- Developed 3D morphology characterization and measured initial thermo-mechanical properties of Ni-Al before and after high-energy ball milling
- Developed WAMR-HPX mini-application for solving the Euler equations

Planned Activities in FY16:

- Continue development of image-based modeling pipeline
- Continue development on dynamic load balancing and fault tolerance; improve network layer, thread scheduler, and lightweight synchronization; optimize AGAS for distributed memory systems and demonstrate with WAMR-HPX
- Calibrate additional material properties and determine parametric uncertainties; validate QOIs (shock speed, final microstructure) for confined impact and assay tests
- Continue DAKOTA immersion with the C-SWARM framework
- Implement boundary tracking algorithm in WAMR; develop PASTA-DDM-HPX using MTL